




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 9/44, 13/14, H04L 12/26, 12/46, 12/66, H04Q 7/00	A1	(11) International Publication Number: WO 00/46669 (43) International Publication Date: 10 August 2000 (10.08.00)																																																					
(21) International Application Number: PCT/US00/02891 (22) International Filing Date: 2 February 2000 (02.02.00) (30) Priority Data: <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">60/118,227</td> <td style="width: 30%;">2 February 1999 (02.02.99)</td> <td style="width: 40%;">US</td> </tr> <tr> <td>09/243,185</td> <td>2 February 1999 (02.02.99)</td> <td>US</td> </tr> <tr> <td>09/243,554</td> <td>2 February 1999 (02.02.99)</td> <td>US</td> </tr> <tr> <td>09/306,678</td> <td>6 May 1999 (06.05.99)</td> <td>US</td> </tr> <tr> <td>09/306,236</td> <td>6 May 1999 (06.05.99)</td> <td>US</td> </tr> <tr> <td>Not furnished</td> <td>28 January 2000 (28.01.00)</td> <td>US</td> </tr> </table> (63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Applications <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">US</td> <td style="width: 30%;">Not furnished (CON)</td> <td style="width: 40%;"></td> </tr> <tr> <td>Filed on</td> <td>28 January 2000 (28.01.00)</td> <td></td> </tr> <tr> <td>US</td> <td>60/118,227 (CON)</td> <td></td> </tr> <tr> <td>Filed on</td> <td>2 February 1999 (02.02.99)</td> <td></td> </tr> <tr> <td>US</td> <td>09/243,185 (CON)</td> <td></td> </tr> <tr> <td>Filed on</td> <td>2 February 1999 (02.02.99)</td> <td></td> </tr> <tr> <td>US</td> <td>09/243,554 (CON)</td> <td></td> </tr> <tr> <td>Filed on</td> <td>2 February 1999 (02.02.99)</td> <td></td> </tr> <tr> <td>US</td> <td>09/306,678 (CON)</td> <td></td> </tr> <tr> <td>Filed on</td> <td>6 May 1999 (06.05.99)</td> <td></td> </tr> <tr> <td>US</td> <td>09/306,236 (CON)</td> <td></td> </tr> <tr> <td>Filed on</td> <td>6 May 1999 (06.05.99)</td> <td></td> </tr> </table>	60/118,227	2 February 1999 (02.02.99)	US	09/243,185	2 February 1999 (02.02.99)	US	09/243,554	2 February 1999 (02.02.99)	US	09/306,678	6 May 1999 (06.05.99)	US	09/306,236	6 May 1999 (06.05.99)	US	Not furnished	28 January 2000 (28.01.00)	US	US	Not furnished (CON)		Filed on	28 January 2000 (28.01.00)		US	60/118,227 (CON)		Filed on	2 February 1999 (02.02.99)		US	09/243,185 (CON)		Filed on	2 February 1999 (02.02.99)		US	09/243,554 (CON)		Filed on	2 February 1999 (02.02.99)		US	09/306,678 (CON)		Filed on	6 May 1999 (06.05.99)		US	09/306,236 (CON)		Filed on	6 May 1999 (06.05.99)		(71) Applicant (for all designated States except US): MENTAT, INC. [US/US]; Suite 315, 1145 Gayley Avenue, Los Angeles, CA 90024 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): TOPOREK, Jerome, D. [US/US]; 343 Rennie Avenue #8, Venice, CA 90291 (US). PALTER, David, C. [US/US]; 10933 Wellworth Avenue #3, Los Angeles, CA 90024 (US). MCCOOEY, Jeremy, A. [US/US]; 1002 Tiverton Avenue, Los Angeles, CA 90024 (US). HASSON, Marc, B. [US/US]; 1722 Goodman Avenue, Redondo Beach, CA 90278 (US). HARTRICK, Timothy, W. [US/US]; 11965 Dorothy Street, Los Angeles, CA 90049 (US). GUYER, Kay, A. [US/US]; 10933 Wellworth Avenue #11, Los Angeles, CA 90024 (US). (74) Agents: BARRETT, Roger, T. et al.; Townsend and Townsend and Crew LLP, Two Embarcadero Center, 8th floor, San Francisco, CA 94111 (US). (81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>
60/118,227	2 February 1999 (02.02.99)	US																																																					
09/243,185	2 February 1999 (02.02.99)	US																																																					
09/243,554	2 February 1999 (02.02.99)	US																																																					
09/306,678	6 May 1999 (06.05.99)	US																																																					
09/306,236	6 May 1999 (06.05.99)	US																																																					
Not furnished	28 January 2000 (28.01.00)	US																																																					
US	Not furnished (CON)																																																						
Filed on	28 January 2000 (28.01.00)																																																						
US	60/118,227 (CON)																																																						
Filed on	2 February 1999 (02.02.99)																																																						
US	09/243,185 (CON)																																																						
Filed on	2 February 1999 (02.02.99)																																																						
US	09/243,554 (CON)																																																						
Filed on	2 February 1999 (02.02.99)																																																						
US	09/306,678 (CON)																																																						
Filed on	6 May 1999 (06.05.99)																																																						
US	09/306,236 (CON)																																																						
Filed on	6 May 1999 (06.05.99)																																																						

(54) Title: INTERNET OVER SATELLITE**(57) Abstract**

A telecommunications apparatus, system and method for providing transport of packetized information over a satellite network, providing a bi-directional flow of information between a client (380) and a server (388) from a first satellite gateway (384) over a connection (386) to a second satellite gateway (386), utilizing protocol translation between TCP protocol and satellite protocol. A telecommunication apparatus, system and method for managing memory for storing information communicated over TCP connections and for controlling information flow in TCP connections are also provided.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNET OVER SATELLITE

5 CROSS-REFERENCES TO RELATED APPLICATIONS

The subject application claims priority from U.S. Provisional Patent Application No. 60/118,227, filed February 2, 1999, in the name of Jerome D. Toporek et. al, titled, "Internet Over Satellite Apparatus," (Attorney Docket Number 16625-001100), the complete disclosure of which is incorporated herein by reference.

10 The subject application also claims priority from the following five commonly-owned co-pending applications, which are hereby incorporated by reference in their entirety for all purposes:

1. U.S. Patent Application Serial No. 09/243,185, filed February 2, 1999, in the name of Jerome D. Toporek et. al, titled, " Internet Over Satellite System,"
15 (Attorney Docket Number 16625-001200);
2. U.S. Patent Application Serial No. 09/243,554, filed February 2, 1999, in the name of Jerome D. Toporek et. al, titled, "Internet Over Satellite Method," (Attorney Docket Number 16625-001300);
3. U.S. Patent Application Serial No. 09/306,678, filed May 6, 1999,
20 in the name of Jerome D. Toporek et. al, titled, "Method and System for Managing Memory in an Internet Over Satellite Connection," (Attorney Docket Number 16625-001210);
4. U.S. Patent Application Serial No. 09/306,236, filed May 6, 1999, in the name of Jerome D. Toporek et. al, titled, "Method and System for Controlling Data
25 Flow in an Internet Over Satellite Connection," (Attorney Docket Number 16625-001220); and
5. U.S. Patent Application Serial No. _____ (Attorney Docket Number 16625-001110), filed January 28, 2000, in the name of Jerome D. Toporek et. al, titled "Internet Over Satellite Apparatus," and which also claims priority
30 from U.S. Provisional Patent Application No. 60/118,227.

BACKGROUND OF THE INVENTION

The present invention relates to telecommunications devices, systems and techniques. More particularly, the present invention provides devices, systems and methods for managing memory and/or transmitting information using a TCP connection over a satellite system for use in a wide area network of computers such as the Internet. The present invention provides a high speed and quality transmission of signals to a distant location over a satellite system using an Xpress Transport Protocol (herein "XTP") protocol, for example. The present invention also provides a controlled memory management in equipment for high speed and quality transmission of signals to a distant location over a satellite system using an XTP protocol, for example. But it will be recognized that the invention has a much wider range of applicability; it can also be applied to a private bridged network, terrestrial wireless network links, and the like.

The Internet is an international "super-network" connecting together millions of individual computer networks and computers. The Internet is generally not a single entity. It is an extremely diffuse and complex system over which no single entity has complete authority or control. Although the Internet is widely known for one of its ways of presenting information through the World Wide Web (herein "Web"), there are many other services currently available based upon the general Internet protocols and infrastructure.

The Web is generally easy to use for people inexperienced with computers. Information on the Web can be presented on "pages" of graphics and text that contain "links" to other pages either within the same set of data files (i.e., Web site) or within data files located on other computer networks. Users often access information on the Web using a "browser" program such as one made by Netscape Communications Corporation of Mountain View, California or Microsoft Corporation of Redmond, Washington. Browser programs process information from Web sites and display the information using graphics, text, sound, and animation. Accordingly, the Web has become a popular medium for advertising goods and services directly to consumers.

The Internet supports many other forms of communication. For example, the Internet allows one-to-one communication via electronic mail, which is commonly known as "e-mail." The Internet also has "bulletin boards" which are used by users to post colorful text and graphics for others to read and respond to. For real time communication, the Internet has "chat rooms" and the like, which allow users to

communicate to each other by way of typed messages. In some cases, the Internet can also be used for voice communication between users. All of these forms of communication are possible, at least in part, by way of an important connection, which allows information to be transmitted from many different servers on the Internet.

5 The Internet is based on the TCP/IP protocol suite. At the network layer, and Internet Protocol, which is known as IP, provides a mechanism to deliver packets addressed to individual computers. The Internet has a plurality of computer systems implementing IP, which are generally interconnected to each other using local area networks and dedicated transmission lines to form a wide area network of computers. IP
10 packets are delivered on a best effort basis and are not guaranteed to arrive at their intended destination, which is known as “unreliable” service.

 For many applications that require reliable data delivery service, the Internet uses a connection mechanism called Transmission Control Protocol, which is known as TCP. TCP has a variety of desirable characteristics that allows it to be suitable
15 for communication over the Internet. For example, TCP considers the data size or best sized segment of data to send from a transmitting server over the communication medium to a receiving server. Additionally, TCP maintains a timer, which waits for the receiving server to send an acknowledgement of the reception of the data segment from the transmitting server. If the timer times out before the acknowledgement is received, TCP
20 resends the data segment, which provides reliability in the communication. TCP also maintains a checksum on its header and data, which monitors any modification in the data in transit. If the data arrives with an invalid checksum, TCP discards the data and does not acknowledge receiving the data. If data segments arrive out of order, TCP can also recombine the segments in order at the receiving server. Furthermore, TCP provides flow
25 control, which only allows a finite amount of data to be sent to a buffer on the receiving server. This prevents a fast server computer from overflowing all the buffers on a slower server computer.

 Although TCP has been well suited for transmitting Internet data between a plurality of servers, which are coupled to each other by way of conventional telephone
30 systems and lines with similar characteristics, TCP has many limitations. For example, TCP is suitable for “short hops” but often slows down communication over extremely long distances which can have significant latency (e.g., transcontinental hops via satellite). Here, the flow control and acknowledgment features take too long to send and

receive, which make communication slow. Additionally, opportunities for gains in efficiency still exist. In particular, TCP has been found to be slow and cumbersome for transmission of Internet data over a satellite network, for example. Further, high speed transmission of Internet data over a satellite network, for example, to a relatively slower recipient, such as modem, and the like, can cause congestion. Additionally, pathological conditions can arise in networking using satellites, that can cause resource exhaustion in systems, which leads to further slowing of TCP connections. Additionally, networking using satellites often encounters significant bit error rates, which leads to further slowing of TCP connections. These and other limitations will be more fully described according to the Figs. below.

From the above, it is seen that a more efficient way of transporting Internet services, and a more efficient way of managing memory in the transmission thereof, over large geographical regions using wireless communication media is highly desirable.

SUMMARY OF THE INVENTION

According to the present invention, a technique for improving TCP performance over a wireless wide area network is provided. In an exemplary embodiment, the present invention provides apparatus, systems and methods for converting information using a TCP connection to an Xpress Transport Protocol (herein “XTP”) connection, for example, which is more suitable for transmission over a wireless network system such as a satellite system.

According to the present invention, techniques for controlling information flow in TCP connections, and for managing memory for storing information communicated over one or more TCP connections over a wireless wide area network are provided. In an exemplary embodiment, the present invention provides methods and systems for controlling the buffering of information and/or the rate that information flows over a TCP connection to an Xpress Transport Protocol (herein “XTP”) connection, for example.

In a specific embodiment, the present invention provides a communication apparatus for transmitting information over a satellite link. The apparatus includes a TCP interface coupled to a satellite gateway interface for providing a bi-directional flow of information, which has data and a header, using a TCP connection. As merely an example, the bi-directional flow of information can be from the Internet or internet-like

network of computers or any network using TCP/IP protocols. The apparatus also includes a processor for converting the information, using the TCP connection, into a satellite protocol for transmission over a satellite link. The satellite protocol has suitable characteristics for transmission of such information over a satellite link. The protocol can include, among others, XTP, XTP-like protocols, and the like.

In an alternative embodiment, the present invention provides a communication apparatus for establishing a plurality of connections for transmission of information using a TCP connection over a satellite link. The apparatus includes a TCP interface for forming a first communication connection between a first client (e.g., user computer) to a first satellite gateway. The apparatus also includes a satellite gateway interface for forming a second communication connection between the first satellite gateway to a second satellite gateway that is over a satellite link. Information describing characteristics (e.g., source and destination IP addresses and port numbers) of the first connection is transmitted to the second satellite gateway. A third communication connection is formed between the second satellite gateway and the destination server. A combination of these connections forms a transparent connection from the original first client to the destination server.

In a specific embodiment, the present invention provides a communication system for transmitting information over a satellite link. The system includes a first gateway having code for providing a bi-directional flow of information, which has data and a header, the first gateway being coupled to a TCP connection. As merely an example, the bi-directional flow of information can be from the Internet or internet-like network of computers or any network using TCP/IP protocols. The system also includes code for converting the information from the TCP protocol into a satellite protocol for transmission over a satellite link. The system also includes a second gateway for receiving the information from the first gateway over a satellite link. The second gateway can include code for converting the information from a satellite protocol, that has suitable characteristics for transmission over a satellite link, to a TCP protocol. The satellite protocol can include, among others, XTP, XTP-like protocols, and the like.

In an alternative embodiment, the present invention provides a communication system capable of establishing a plurality of connections for transmission of information using a TCP connection over a satellite link. The system includes code for intercepting a first communication connection between a first client (e.g., user computer)

to a first satellite gateway. The system also includes code for forming a second communication connection between the first satellite gateway to a second satellite gateway that is over a satellite link. Information describing characteristics (e.g., source and destination IP addresses and port numbers) of the first connection is transmitted to the second satellite gateway. A third communication connection is formed between the second satellite gateway and the destination server. A combination of these connections forms a transparent connection from the original first client to the destination server.

In a specific embodiment, the present invention provides a communication method for transmitting information over a satellite link. The method includes providing a bi-directional flow of information, which has data and a header, using a TCP connection. As merely an example, the bi-directional flow of information can be from the Internet or internet-like network of computers or any network using TCP/IP protocols. The method also includes converting the information, using the TCP connection, into a satellite protocol for transmission over a satellite link. The satellite protocol has suitable characteristics for transmission of such information over a satellite link. The protocol can include, among others, XTP, XTP-like protocols, and the like.

In an alternative embodiment, the present invention provides a communication method for establishing a plurality of connections for transmission of information using a TCP connection over a satellite link. The method includes intercepting a first communication connection between a first client (e.g., user computer) to a first satellite gateway. The method also includes forming a second communication connection between the first satellite gateway to a second satellite gateway that is over a satellite link. Information describing characteristics (e.g., source and destination IP addresses and port numbers) of the first connection is transmitted to the second satellite gateway. A third communication connection is formed between the second satellite gateway and the destination server. A combination of these connections forms a transparent connection from the original first client to the destination server.

In a specific embodiment, the present invention provides a method for managing memory for buffering information communicated over an internet connection established across a satellite link. The method can be operable on a large variety of systems, such as a system comprising one or more clients, one or more servers, a first gateway connected to at least one client, a second gateway connected to at least one server, and a wireless telecommunications link between these gateways. The method

includes intercepting a connection attempt with a server initiated by a client. The client will attempt a connection that will have a first characteristic data rate. A connection can be established between a first gateway and a second gateway over the telecommunications link, that can be a satellite link or the like. This connection will have a second characteristic data rate. The first gateway can include a buffer for storing information received over the telecommunications link from the second gateway and a buffer for storing information to be sent to the client. The method can perform a step of determining a state in the receive buffer. The state can indicate whether the characteristic data rate in the telecommunications link is substantially greater than the characteristic data rate for the client. A step of setting a receive window size for the client at the first gateway to reduce an amount of data received over the telecommunications link for the client can also be part of the method. The combination of these steps can provide a method for managing memory for buffering information coming from the telecommunications link in the first gateway.

In another aspect of the present invention, a system for controlling the flow of information over a satellite is provided. The system can include at least one client, selected from a plurality of potential clients and at least one server, selected from a plurality of potential servers. The system can also include a first gateway, connected to the client by a first telecommunications link, and a second gateway, connected to the server by a second telecommunications link. A third telecommunications link connecting the first gateway to the second gateway can also be part of the system. The system is operative to intercept a connection attempt with the server initiated by the client at a first characteristic data rate. The system can then establish a connection between the first gateway and the second gateway over the third telecommunications link. This connection will have a second characteristic data rate. The first buffer stores information received over the third telecommunications link and the second buffer stores information to be sent to the client. The system can determine a state in the first buffer. The state indicates that the second characteristic data rate in the third telecommunications link is greater than the first characteristic data rate for the client. The system can also set a receive window size for the client at the first gateway to reduce an amount of data received over the third telecommunications link for the client.

In a yet further aspect according to the present invention, a computer program product for controlling the flow of information over a wireless data link is

provided. The data link can comprise one or more connections. The computer program product comprises a computer readable storage medium for holding a plurality of code. Code for intercepting a connection attempt with a server initiated by a client at a first characteristic data rate is included in the program product. Further, code for establishing
5 a connection between a first gateway and a second gateway over a telecommunications link at a second characteristic data rate is also included. A first buffer in the first gateway stores information received over the telecommunications link and a second buffer stores information to be sent to the client. The product can also include code for determining a state in the first buffer. The state indicates a condition where the characteristic data rate
10 of the telecommunications link is substantially greater than that of the client. If this condition is detected, then the product invokes code for setting a receive window size for the client at the first gateway to reduce an amount of data received over the telecommunications link for the client.

In a specific embodiment, the present invention provides a method for
15 controlling the flow of information over an internet connection established across a satellite link. The method can be operable on a large variety of systems, such as a system comprising one or more clients, one or more servers, a first gateway connected to at least one client, a second gateway connected to at least one server, and a wireless telecommunications link between these gateways. The method includes intercepting a
20 connection attempt with a server initiated by a client. A connection can be established between a first gateway and a second gateway over the telecommunications link, that can be a satellite link or the like. The method can include a step of determining a round trip time for data to travel from the client to the server over the connection. A data rate for the connection can also be determined. From the round trip time and the data rate, a
25 bandwidth delay product can be determined by the method. The method can also determine a flow control limit for the connection from the bandwidth delay product. A step of limiting data transfer over the connection based upon the flow control limit can be part of the method. The combination of such steps as these can provide a method for controlling the flow of TCP data over networks spanning large geographical regions using
30 a wireless communication media, and the like.

In another embodiment, the present invention can select from among connections, ones that can benefit available system resources by being flow controlled. Some embodiments can screen for flow control those connections that are not initializing,

or are transferring data in blocks below a threshold value, or above a threshold value and the like. In many embodiments, a bandwidth delay product for the connection can be determined by multiplying round trip time and data rate. In some embodiments, a flow control limit for a connection can be determined by multiplying bandwidth delay product
5 by a scaling factor and dividing by a number of connections for the link.

In another aspect of the present invention, a system for controlling the flow of information over a satellite is provided. The system can include at least one client, selected from a plurality of potential clients and at least one server, selected from a plurality of potential servers. The system can also include a first gateway, connected to
10 the client by a first telecommunications link, and a second gateway, connected to the server by a second telecommunications link. A third telecommunications link connecting the first gateway to the second gateway can also be part of the system. The system is operative to determine a round trip time for data to traverse from the client over the first telecommunications link, the second telecommunications link and the third
15 telecommunications link to the server over a connection. The system can determine a data rate for the connection. Further, the system can determine a bandwidth delay product from the round trip time and the data rate. The system can determine a flow control limit for the connection from the bandwidth delay product and limit data transfer over the connection based upon the flow control limit. Limiting data transfer in such a
20 manner can interface relatively fast data transfer on the telecommunications links to a relatively slower data transfer rate in the client, for example.

In a yet further aspect according to the present invention, a computer program product for controlling the flow of information over a wireless data link is provided. The data link can comprise one or more connections. The computer program
25 product comprises a computer readable storage medium for holding a plurality of code. Code for determining a round trip time for data to traverse the wireless data link over a connection can be part of the product. Further, the computer program product can include code for determining a data rate for the connection. Code for determining a bandwidth delay product from the round trip time and the data rate can also be included. Code for
30 determining a flow control limit for the connection from the bandwidth delay product and code for limiting data transfer over the connection based upon the flow control limit can also be part of the computer program product.

In another aspect of the present invention, techniques for matching a data rate for information flowing through a predetermined point in a network, such as a satellite, to the data rate capabilities of a remote point on the network are provided. In one embodiment, the present invention provides rate control to maintain a suitable or predetermined data rate over the satellite device. The rate control provides a plurality of queues for buffering incoming data. Incoming data are checked against a predetermined length. Data in excess of the predetermined length are stored in the queues. At predetermined intervals, information is taken from one or more queues and transmitted along the outgoing data path. This processing can enable systems according to the present invention to provide for data rate control across a particular point of interest along the network, such as a satellite link.

In another aspect of the present invention, a method of converting a communications system for providing transport of packetized information in a TCP protocol between a client and a server to a system for providing transport of packetized information in a satellite protocol between the client and the server includes a step of installing a first gateway operatively disposed to intercept connections from the client to the server. The first gateway is further able to convert the packetized information from TCP protocol to satellite protocol. A step of installing a second gateway capable of establishing a connection with the server and further operatively disposed to convert packetized information from satellite protocol to TCP protocol is also included in the method. The first gateway and the second gateway are able to establish a connection over a common telecommunications link. The combination of these elements can provide a method for converting a system for TCP telecommunications to satellite protocol communications.

Numerous benefits are achieved by way of the present invention over conventional techniques. In a specific embodiment, the present invention provides a higher throughput than conventional systems. Additionally, the present invention provides for a way to transparently improve TCP performance over a satellite network. In other embodiments, the present invention provides for a novel protocol which is suitable for long latency, high loss, asymmetric bandwidth conditions, which are typical of satellite communications.

In a specific embodiment, the techniques according to the present invention can quickly resize the receive window to achieve a match between the data rate

of a sender and the data rate of a receiver of a TCP connection over a satellite link. In some embodiments, the techniques of the present invention can be applied to each satellite connection separately, so that one slow receiver will not degrade performance of other connections over the same link. In some embodiments, the present invention provides for a management of buffer memory in a novel protocol which is suitable for long latency, high loss, asymmetric bandwidth conditions, which are typical of satellite communications. In a specific embodiment, the present invention provides a more even distribution of network resources than conventional systems. Many embodiments according to the present invention provide for sharing a network/satellite link between connections. In many embodiments, it is less likely that a burst of data from one particular connection will be queued for transmission over the link in front of data for other connections. In other embodiments, the present invention provides for data flow control in a long latency, high loss, asymmetric bandwidth communications medium, which is typical of satellite communications.

The present invention also provides for relatively easy implementation into, or interface with, a conventional satellite system. Depending upon the embodiment, one or more of these benefits can be present. These and other advantages or benefits are described throughout the present specification and are described more particularly below.

These and other embodiments of the present invention are described in more detail in conjunction with the text below and attached Figs.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a simplified diagram of a satellite system according to an embodiment of the present invention;

Figs. 2 and 2A are simplified diagrams of system architectures according to embodiments of the present invention;

Figs. 3A-3J are simplified diagrams of process diagrams according to embodiments of the present invention;

Figs. 4-6 are simplified diagrams of experimental data according to embodiments of the present invention; and

Fig. 7 is a simplified diagram of a gateway interface according to an embodiment of the present invention.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

Before discussing the various embodiments, it may assist the reader to more fully understand the limitations of using TCP over a satellite network. Some of these limitations are described more fully below.

Congestion Avoidance: In order to avoid the possibility of congestive network meltdown, TCP generally assumes that all data loss is caused by congestion and responds by reducing the transmission rate. Over satellite links, however, TCP often misinterprets the long round trip time and bit errors as congestion and responds inappropriately. Similarly, the TCP "Slow Start" algorithm, which over the terrestrial infrastructure prevents new connections from flooding an already congested network, over satellites forces an excessively long ramp-up for each new connection. While these congestion avoidance mechanisms are important in routed environments, they are often ill-suited to satellite links.

Data Acknowledgments: The simple, heuristic data acknowledgment scheme used by TCP generally does not adapt well to long latency or highly asymmetric bandwidth conditions. To provide reliable data transmission, the TCP receiver often constantly sends acknowledgments for the data received back the sender. The sender generally does not assume any data is lost or corrupted until a multiple of the round trip time has passed without receiving an acknowledgment. If a packet is lost or corrupted, TCP can retransmit all of the data starting from the first missing packet. The algorithm may not respond well over satellite networks where the round trip time is long and error rates are high. Further, the constant stream of acknowledgments frequently wastes precious back channel bandwidth. If the back channel bandwidth is small, the return of the acknowledgments to the sender can become the system bottleneck in some cases.

Window Size: TCP utilizes a sliding window mechanism to limit the amount of data in flight. When the window becomes substantially full, the sender stops transmitting data until it receives new acknowledgments from the destination server. Over satellite networks, where acknowledgments are slow to return, the TCP window size generally sets a hard limit on the maximum data throughput rate. The minimum window size often needed to fully utilize an error-free link, known as the "bandwidth-delay product," is 100 Kbytes for a T1 satellite link and 675 Kbytes for a 10 MBps link, for example. Bit errors can increase the required window size further. However, most implementations of TCP/IP are often limited to a maximum window size of 64 KB and many common operating systems use a default window size of only 8 KB, imposing a maximum throughput rate over a satellite link of only 128 Kbps per connection, regardless of the bandwidth of the data pipe.

According to the present invention, a technique for providing a TCP connection over a wide area wireless network is provided. In an exemplary embodiment, the present invention provides apparatus, systems and methods for converting information using a TCP connection to an XTP connection, which is more suitable for transmission over a wireless network system such as a satellite system.

Further, according to the present invention, techniques for controlling information flow in TCP connections, and for managing memory for one or more TCP connections over a wireless wide area network are provided. In an exemplary embodiment, the present invention provides methods and systems for controlling the flow of information and/or memory management for buffering information transmitted over a TCP connection to an XTP connection, which is more suitable for transmission over a wireless network system such as a satellite system. Details of the present invention are described throughout the present specification and more particularly below.

Fig. 1 is a simplified diagram of a satellite system 100 according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Among other features, the system 100 includes a satellite network system, which has satellite(s) 101. The satellite 101 receives

and transmits information between a plurality of ground stations 107, 108 or satellite dishes. Each satellite ground station includes a satellite dish, and a satellite modem 109A, 109B or other form of receiver such as a VSAT indoor unit, which is coupled, directly or indirectly to a satellite gateway 111A, 111B.

5 Satellite antenna 108 receives and transmits signals 103 through satellite 101. Satellite antenna 108 connects to satellite gateway 111B that couples to a wide area network such as the Internet 129 or an internet, which is coupled to a server 131. The gateway 111B also couples through satellite modem 109B. The server and the Internet communication occur in a common protocol such as TCP/IP, UDP/IP or the like. The
10 satellite gateway will be described in more detail below.

 Satellite antenna 107 receives and transmits signals 105 through satellite 101. Satellite antenna 107 connects to satellite gateway 111A that couples to a local area network such as Ethernet, Token Ring, FDDI, or others 121, which is coupled to computer terminals 123 or clients. Here, satellite gateway couples to laptop 117, which is
15 coupled through modem 119. The client, the laptop, and the local area network use a common connection format such as TCP/IP, IPX/SPX, or the like. The satellite gateway will be described in more detail below.

 In a specific embodiment, satellite gateway 111A intercepts a TCP connection from a client and converts data 105 to a satellite protocol for transmission
20 over satellite 101. The gateway 111A also couples through satellite modem 109A, which transmits data to satellite 101. The satellite gateway 111B on the opposite side of the satellite link translates the data in the satellite protocol back to TCP for communications with the server 131. In a specific embodiment, the present invention provides improved performance over conventional techniques. Additionally, techniques of the present
25 invention remain substantially transparent to an end user and is fully compatible with the Internet or an internet infrastructure. In many, if not all aspects, no substantial changes are required to the client or server. In preferred embodiments, applications continue to function without substantial hardware and/or software modifications.

 Although the above has been described in terms of a specific networking
30 diagram, it is possible to implement the present satellite gateway in other configurations. For example, the present satellite gateway can be one "hub" or central gateway that serves a plurality of remote gateways. Each of the remote gateways is connected to the central gateway to create an individual satellite link.

Other common network topologies can also be used. Further, in some embodiments, different telecommunications links can be used to carry forward and backward paths of a connection. For example, a satellite link can be used to carry data on the forward path, while telephone line can be used for the backward path. Other types of telecommunications hardware may be substituted for the example hardware described here without departing from the scope of the present invention.

Additionally, the present satellite gateway is generally described as a stand alone unit. It will be recognized, however, that the present gateway can be implemented or integrated into a client machine. For example, the present gateway can be implemented on a personal computer using a satellite card, which can be inserted into a WindowsTM 98 machine. The present gateway can also be implemented on other operating systems such as Windows NT, MacOS, and Linux. Of course, the exact manner of implementation will depend upon the application. Additionally, the present satellite gateway may be integrated into a standalone satellite modem, VSAT hardware, router or other network devices.

I. PROTOCOL DESIGN

In a specific embodiment, the present invention includes a novel satellite protocol, which provides improved throughput for satellite networks. The present protocol can be designed to respond efficiently to typical satellite latency, bit errors, and asymmetric bandwidth conditions. The present protocol also can take advantage of optimizations possible on a point-to-point link with fixed bandwidth. For further information regarding satellite protocols, such as XTP, reference may be had to Tim Strayer, "A Brief Introduction to the Xpress Transport Protocol," which is incorporated herein by reference in its entirety for all purposes. These and other features of the present satellite protocol are described in more detail below.

The present protocol utilizes an efficient selective retransmission algorithm for the acknowledgment of data. Because the hop over the satellite is a point-to-point link with no intermediate routing, any gaps in the packet sequence can be assumed to be data loss due to corruption. The receiving satellite gateway requests retransmission immediately upon detection of the missing data from the transmitting satellite gateway.

The present protocol is substantially free from the frequent acknowledgment features of conventional TCP. In some embodiments, the present protocol does not use acknowledgments as the primary means of identifying lost data. In these embodiments, the present protocol needs only infrequent acknowledgments to confirm data arrival and clear buffers. (Conventional TCP sends a constant stream of acknowledgments over the reverse channel.) The present protocol reduces back channel usage by 70% or more, thereby increasing the performance of networks where limited back channel bandwidth is the system bottleneck.

The present protocol offers adequately large window sizes for transmission of data between the satellite gateways. Because the bandwidth-delay product over the satellite between the satellite gateways is much larger than that from the satellite gateway to the end node, the large present protocol window enables the bandwidth-delay product to window size ratio to remain relatively constant. The present gateway becomes a buffer for the network, allowing high throughput independent of the window size of the clients and servers.

The present protocol generally does not use unnecessary congestion avoidance mechanisms for the hop over the satellite between the satellite gateways. The present system, however, continues to use "Slow Start" and all other standard TCP congestion avoidance algorithms for the terrestrial connections between the present gateways and the end nodes. The present protocol also uses a streamlined handshake mechanism to reduce the number of round-trip times required for connection set-up.

The present apparatus, systems and methods can run over IP for compatibility with routed networks. Alternatively, it can run directly over the link layer for substantially improved performance in most instances. Depending upon the embodiment one or more of these features can be available.

The above has generally been described in terms of desirable features of the present satellite protocol. It would be recognized that many other types of features can be available. Additionally, the present invention does not necessarily require each of the above features. Any combination can be used depending upon the application. One of ordinary skill in the art would recognize other variations, modifications, and alternatives.

II. XPRESS TRANSPORT PROTOCOL

In a preferred embodiment, the present invention provides an "Xpress Transport Protocol," which is commonly known as XTP. XTP provides orthogonal protocol functions for separation of implementation from policy, separation of rate and flow control, explicit first-class support for multicast, and data delivery service independence. XTP has a variety of suitable characteristics for transmission of data over a satellite link, which are described more fully below.

In a specific embodiment, the present protocol includes a set of mechanisms whose functionality is orthogonal. The most notable effect of this is that XTP separates communication mechanism (e.g., data-gram, virtual circuit, transaction, etc) from an error control policy employed (fully reliable through uncorrected). Further, flow and rate control as well as error control can be tailored to the communication at hand. If desired, any set of these control procedures can be turned off.

The present protocol also provides for separation of rate and flow control. In general, flow control operates on end-to-end buffer space. Rate control is a producer/consumer concept that considers processor speed and congestion. TCP does not provide rate control, and combats congestion with slow-start and other heuristics. XTP provides mechanisms for shaping rate control and flow control independently.

In a specific embodiment, the present protocol provides explicit multicast support. Here, multicast is not an afterthought in XTP. Each mechanism used for uni-cast communications is available for multicast use as well.

The present protocol also has data delivery service independence. In particular, XTP is a transport protocol, yet with the advent of switched networks rather than routed networks, a traditional network layer service may not be appropriate in every instance. XTP generally requires that the underlying data delivery service provides framing and delivery of packets from one XTP-equipped host to another. This could be raw MAC or IP or AAL5. XTP also employs parametric addressing, allowing packets to be addressed with any one of several standard addressing formats. Other features of XTP include, among others, implicit fast connection setup for virtual circuit paradigm; key-based addressing lookups; message priority and scheduling; support for encapsulated and convergence protocols; selective retransmission and acknowledgment; and fixed-size 64-bit aligned frame design.

XTP defines the mechanisms necessary for delivering user data from one end system to one or more other end systems. Each XTP implementation maintains the state of each of its communications. Well-defined packet structures, containing user data or control information, are exchanged in order to effect the user data transfer. The control information is used to provide the requested layer of correctness and to assist in making the transfer efficient. Assurance of correctness is done via error control algorithms and maintenance of a connection state machine. Flow and rate control algorithms, certain protocol modes, and traffic shaping information are used to provide the requested quality of service as efficiently as possible.

The collection of information comprising the XTP state at an end system is called a context. This information represents one instance of an active communication between two (or more) XTP endpoints. A context should be created, or instantiated, before sending or receiving XTP packets. There may be many active contexts at an end system, one for each active conversation or message.

Each context manages both an outgoing data stream and an incoming data stream. A data stream is an arbitrary length string of sequenced bytes, where each byte is represented by a sequence number. The aggregate of a pair of active contexts and the data streams between them is called an association.

A context at an end system is initially in a quiescent state. A user in need of communication services requests that the context be placed into the listening state. The context now listens for an appropriate FIRST packet. The FIRST packet is the initial packet of an association. It contains explicit addressing information. The user should provide all of the necessary information for XTP to match an incoming FIRST packet with the listening context.

At another end system a user requests communication service from XTP. Since this user will initiate the association, the context moves from a quiescent state to an active state directly. The active context constructs a FIRST packet, complete with explicit addressing information from the user. The FIRST packet is sent via the underlying data delivery service.

When the FIRST packet is received at the first host's XTP implementation, the address is compared against all listening contexts. If a match is found, the listening context moves to the active state. From this point forward an association is established, and communication can be completely symmetric since there are two data streams, one in

each direction, in an association. Also, no other packet during the lifetime of the association will carry explicit addressing information. Rather, a unique "key" is carried in each packet that maps the packet to the appropriate context.

Once all data from one user has been sent, that data stream from that user's context can be closed. Sentinels in the form of options bits in a packet are exchanged to gracefully close the connection. Other forms of less graceful closings are possible by abbreviating this exchange. When both users are done, and both data streams closed, the contexts move into the inactive state. One of the contexts will send a sentinel that causes the association to dissolve. At this point, both contexts return to the quiescent state.

Generally all of XTP's packet types use a common header structure. All of the information necessary to steer the packet's payload to the proper point of processing is carried in the header. Much of how an XTP context operates is controlled by a set of bit flags that are concentrated in one field in the packet header. Fifteen flags are defined, including bit flags to control connection shutdown, bit flags to control the acknowledgment policy, and bit flags that are markers in the data stream. The remaining bit flags control the end-to-end operating modes. Examples include enabling or disabling error control or flow control, or enabling multicast mode.

XTP flow control is based on 64-bit sequence numbers and a 64-bit sliding window. XTP also provides rate control whereby an end system or intermediate system can specify the maximum bandwidth and burst size it will accept on a connection. A Traffic Segment provides a means for specifying the shape of the traffic so that both end systems and intermediate systems can manage their resources and facilitate service quality guarantees.

Error control in XTP incorporates positive and, when appropriate, negative acknowledgment to effect retransmission of missing or damaged data packets. Retransmission may be either go-back-N or selective. The retransmission algorithms are conservative: only data that is shown to be missing via control messages may be transmitted. This avoids spurious and possible congestion-causing retransmissions. The error control algorithm, while basically conservative, can also be aggressive: systems, techniques and methods for a quick-acting error notification are provided.

XTP also specifies techniques for extending error control to a multicast environment. The error control algorithm in multicast is identical to the uni-cast algorithm, although additional sophistication is required to manage state variables and

achieve continuous streaming. Accordingly, XTP is a preferred satellite protocol according to the present invention.

Although the above has generally been described in using an XTP protocol, it will be understood that other types of protocols can be used. For example, the protocol can be a modified TCP, a modified XTP, and others. Additionally, the protocol can be any that is suitable for transmission over a satellite link. Additionally, one of ordinary skill in the art would recognize other variations, modifications, and alternatives.

Fig. 2 is a simplified diagram of system architecture 200 according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. The system architecture 200 includes at least a client 201, which is coupled to a satellite gateway 203, which transmits and receives information via satellite link 209 from a satellite gateway 205. Satellite gateway 205 couples to server 207. Satellite gateway and server couple to each other through the Internet or an internet-like network of computers.

Client 201 can be represented in multiple layers, including application and physical layers. The layers may include a browser 211. The browser 211 allows a user to communicate information from an application layer to a physical layer for transmission to a gateway. The browser 211 is generally in the application layer, which provides the information. For example, the browser can be one of suitable design made by a company called Netscape Communications Corporation of Mountain View, California or Microsoft Corporation of Redmond, Washington or others. In addition to browsers, other TCP applications, including FTP file transfers, may also be used to communicate between clients and servers.

The information goes through the transport layer (e.g., TCP) 213 and then through the IP layer 215, which is the networking layer. The transport layer provides a flow of data between the client and the gateway. The IP layer handles movement of data comprising packets between the client and the gateway, or other network. The information is sent through the physical layer, which includes a driver 217. The driver 217 transmits the information to the gateway 203 through hardware 219 connected to gateway 203 via a telecommunications link 221, which can be a wire, cable, fiber optic, or other physical communication medium. Alternatively, the driver 217 receives information from the gateway 203 via link 221 and hardware 219. The driver can be a

network operating system with a network interface card in the client computer, for example.

From the physical connection 221, the information traverses to the gateway 203. The gateway has a physical layer 223, which interacts with driver 225. The gateway also includes a networking layer 227 and a transport layer 229, which is coupled to a translation module 231. The networking layer 227 determines whether the information can be routed over the satellite protocol. If so, the data is passed up to the transport layer 229. If not, the data is immediately forwarded out to the rate control module 234 for delivery to the satellite link 239 in non-translated form. The translation module 231 converts the information into a satellite protocol 233, which is suitable for use over a satellite link. A rate control module 234 determines whether the information can be passed immediately to the satellite connection or be queued for later delivery. The satellite connection is coupled to a physical layer 237 by a driver 235, which transmits information to and from satellite 209. The information traverses to and from the satellite in a wireless medium, 239, 241.

Information is received by the satellite gateway 205, which includes multiple layers, physical and others. The physical layer 243 couples to driver 245. The networking and transport layer include a satellite protocol 247 and a rate control 244. The network and transport layers connect to a session layer which comprises translation module 249. The translation module converts the information using the satellite protocol back to TCP/IP. The information traverses through the transport layer (i.e., TCP) 251 and the networking layer (i.e., IP) 253. Information from the satellite gateway 205 enters a physical layer 257 through driver 255. The driver transmits the information to a server 207 over telecommunications link 259. Driver 255 can also receive information from server 207 in the backward path.

From the telecommunications link 259, the information is sent to driver 263 in the server 207 via physical layer 261. The information traverses from the driver to a networking layer, which includes IP 265. The information also traverses from the networking layer to a transport layer, which includes TCP 267. The information is delivered to a Web server application 269, for example. Between the server 207 and the satellite gateway 205 is the Internet or another IP network, depending upon the application.

In a specific embodiment, information can flow through a gateway, such as gateway 203, 205 at the network layer, bypassing the transport and application layers altogether. For example, in gateway 203, information can enter from client 201 via telecommunications link 221. Physical layer 223 passes the information to driver 225. In turn, driver 225 passes the information to network layer 227. Network layer 227 can route the information to its destination using IP. The information then flows from network layer 227 to driver 235. Driver 235 interacts with physical layer 237 to pass the information along to satellite 209 via telecommunications link 239.

In a specific embodiment, the translation module converts information using a TCP connection to a suitable connection for transmission over a satellite system. The suitable connection for the satellite is generally resilient to transmission over a wireless network such as a satellite network, e.g., Orion Worldcast™ or PanAmSat™ SpotBytes™ services, and the like. The connection should also be suitable for long latency, high loss, and asymmetric bandwidth conditions. The connection should also be transparent to the end user at the client or server location. The long latency is typically about 200 milliseconds to about 700 milliseconds, per satellite hop but is not limited to such values.

In a specific embodiment, the present translation module converts information using a TCP connection into an XTP, modified TCP or XTP-like protocol for transmission over a satellite link, which is illustrated by Fig. 2A, for example. Here, the TCP module receives information 350 including the TCP connection. The information includes data 333, a TCP header 335, and an IP header 337. The TCP module removes the TCP header from the information such that the information 351 includes substantially all data 333 and passes the data, along with certain TCP header information, to the translation module. The translation module passes the data or portions of data to the satellite protocol module where a satellite protocol header 339 is prepended onto the data. The header is an XTP header or like. The information 352 now includes the XTP header, which is suitable for transmission over the satellite link. Depending on the specific implementation, the XTP header and data may also be prepended with an IP header for transmission over the satellite link. In particular, the information is transferred to the physical layer and the driver. The driver and physical layer send the information to a receiving satellite gateway, which may prepend additional headers.

From the satellite link, the information including the XTP header enters a receiving XTP module. The receiving XTP or satellite protocol module removes the XTP header from the information, leaving data 353. The receiving satellite protocol module passes the data to the translation module, which can be a routing module used to route data to the proper protocol. The translation module then passes the data to the TCP module. The receiving TCP module attaches a TCP header 343 and IP header 341 onto the data to form information that is now ready for transmission over a network to a receiving server destination. The information using the TCP connection traverses through a network to the destination server.

In other embodiments, the translation module can also convert a portion of the information including the data and TCP/IP headers to information using the XTP protocol. Alternatively, the translation module can convert more than one segment of information including multiple TCP/IP headers into a single piece of information including the XTP header for transmission over a satellite link. Depending upon the application, the translation module can convert the information including the TCP connection into any one or any combination of the above. Further details of these techniques, and others, are described in more detail below.

Figs. 3A and 3B are simplified diagrams of process diagrams according to embodiments of the present invention. These diagrams are merely illustrations and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. In a specific embodiment, the present invention uses a connection process 300, which is illustrated by Fig. 3A. The connection process uses a plurality of separate connections using a "handshaking routine" in the satellite system to provide a transparent TCP connection to an end user. The satellite system can be similar to the one described herein, but is not limited. The present process begins at start, step 301. In a specific embodiment, a plurality of connections are provided. In particular, the present process provides a first connection, step 303. The first connection is a TCP connection between a client and a client side gateway, which interfaces to a satellite link. The present process provides a second connection, step 305, which provides a connection between the client side gateway and a destination gateway over a wireless (e.g., satellite) link. The connection between the wireless link can be any suitable connection such as XTP or XTP-like connection. The present process provides a third connection, step 307, which forms a TCP connection between the destination

gateway to the destination server, which can traverse through the Internet or an internet. Once these connections have been made they are maintained, step 309. The present process ends at stop 311, which terminates all three connections.

5 In a specific embodiment, the present process occurs by way of a selected sequence of steps. Here, TCP SYN packets are intercepted transparently by a satellite gateway. The satellite gateway establishes a new XTP connection across the satellite link to the other satellite gateway. Information, including IP addresses and port numbers, about the requesting client and the destination server is transferred across the satellite link to the other satellite gateway. The destination gateway then sets up a TCP connection
10 with the destination server, using the client's addressing information so that the server sees the TCP connection as being the same as if the client itself had established the connection. The satellite gateway makes routing decisions based upon source and destination addresses in the IP header of packets entering the gateway in order to determine whether the packets should be routed over the satellite telecommunications
15 link. Once the TCP connection to the server is established, the destination gateway communicates back to the sending gateway, which then acknowledges the connection with the client.

In one or more embodiments, this sequence of steps has advantages. In particular, the client does not believe the connection is established until the server has
20 accepted this connection, which tends to reduce or even eliminate false indications of successful connection establishment. Additionally, both the client and server see the connection as happening immediately between the two of them, without detecting that the satellite connection is happening in between the client and the server. That is, the present invention substantially preserves "end-to-end semantics" of the connection. In other
25 embodiments, the connectivity is symmetric. Here, "clients" may exist on either side of the satellite link, and "servers" may exist on either side. Systems on either side may request connections to systems on the other side, and the satellite gateways will intercept the connections.

Fig. 3B illustrates a simplified flowchart of a novel routing process 320 in
30 a specific embodiment according to the present invention that begins with a starting step 321. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. In a step 323, a client host sends a TCP packet destined for a particular

IP address to a satellite gateway. In an example, the satellite gateway can be similar to the one 111A described above, but can also be others. In a decisional step 325, the satellite gateway receives the packet and determines whether the packet is to be carried over an XTP connection prior to transmission. An example of an XTP protocol has been described herein, but should not limit the scope of the claims. Any suitable satellite protocol can be used according to certain embodiments of the present invention. If the destination address of the TCP packet is configured to be converted to the XTP protocol format, then the packet is converted into the XTP protocol in a step 331. Otherwise, the packet is transmitted in its TCP format via an alternative branch to step 327. By way of the present decisional process, the present invention is useful for sites where a single satellite gateway communicates with multiple remote sites, some of which have gateways that are compatible with XTP protocol while other sites do not contain corresponding gateways. Additionally, the present invention allows specific addresses to be configured for protocol translation and other addresses to be configured for no translation, based on policy decisions of the network administrator.

In a specific embodiment, the present invention also provides a rate control step (step 327). In a step 327, processing is performed to maintain a suitable or predetermined data rate over the satellite device. The processing of step 327 may require buffering of packets in order to avoid overloading the satellite link, such as 105 and 103 in Fig. 1. This can be accomplished using the process depicted by Figs. 3C-3D, but is not limited to this process. Then, in a step 329, the packet is transmitted along a connection transmitted from the sending satellite gateway across the satellite link. Then, in a step 333, the packet is converted back to TCP, if it was converted to XTP protocol in step 331. Then, in a step 335, the packet is routed to its remote destination. The present invention is especially useful for sites where a single satellite gateway communicates with multiple remote gateways, some of which are compatible with XTP protocol while others of the remote gateways are not. Some configurations may have multiple remote sites per satellite gateway. Some of these remote sites may have a satellite gateway installed, while others do not. Thus, there may be no satellite gateway on the remote side, just a IP routing configuration. The process completes routing (step 335) when the information enters a destination server 131, which transfers the information to a destination client. The process stops (step 337) when the connection is terminated.

In a particular embodiment according to the present invention, TCP connections at times pass a piece of information in the protocol header which must be delivered as soon as possible to the destination side of the connection. This piece of information is known as the “urgent pointer”. In select embodiments, the TCP implementation that is part of the client-side satellite gateway, such as for example satellite gateway 203 of Fig. 2, can extract this urgent pointer from the TCP header. However, this operation is not required by the system of Fig. 2 and is not intended to limit the scope of the claims. A TCP module, such as TCP module 229 of Fig. 2, for example, can extract the urgent pointer and pass it to a satellite gateway translation module, which can be satellite gateway translation module 231 of Fig. 2, or an equivalent. The urgent pointer is passed down to a satellite protocol module, such as satellite protocol module 233 of Fig. 2, or an equivalent. The urgent pointer can then be carried in the satellite protocol header to a satellite gateway, such as satellite gateway 205 of Fig. 2, or any other receiving satellite gateway. At the receiving satellite gateway, the urgent pointer can be extracted by a satellite protocol module, such as satellite protocol module 247 of Fig. 2, and passed up to a translation module, such as translation module 249, or another satellite translation module, which can deliver it to a TCP module, such as TCP module 251. The TCP module then incorporates the urgent pointer in its target field in the TCP header for immediate delivery to an end server, such as end server 207. This header for delivery to the server refers to the point in the data stream that has not yet arrived at the second satellite gateway machine. In some embodiments, appropriate urgent pointer processing can alleviate malfunctions.

Fig. 3C illustrates a simplified flowchart of a data rate control process such as data rate control step 327 in Fig. 3B in a specific embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3C shows a first decisional step 341 of determining whether queues are empty when an incoming packet arrives. In a presently preferable embodiment, there are two queues, one for queuing high priority (HIPRI) traffic and one for queuing “normal” traffic. A person of ordinary skill in the art can appreciate that the number of queues can be extended or reduced by straightforward extensions of the embodiments according to the present invention. If the queues are empty, then

processing continues with a decisional step 343. Otherwise, processing continues with a decisional step 345.

Decisional step 343 determines whether the current system clock tick count is the same as a stored clock tick count. If the system clock tick and the stored clock tick are the same, then processing continues with decisional step 349, which determines whether a byte counter and a length of an incoming packet are greater than a predetermined maximum, based on the desired transmission rate, denoted here as "MAX". If the maximum would not be exceeded, then processing continues with step 355, in which the packet is sent and the byte counter is incremented by the packet length. Otherwise, if the maximum would be exceeded in step 349, then processing continues with a step 357, in which the difference between MAX and the current value of the byte counter is stored as "EXTRA", and a timer can be started. Processing then continues at step 345. Otherwise, if step 343 determines that the clock tick and stored clock tick are different, processing continues with a step 347 in which a byte counter is cleared and the stored clock tick is updated. Processing then continues with a step 355, as described above.

If decisional step 341 determines that at least one of the queues is not empty, or if the packet length would have increased the byte counter beyond MAX in step 349, then in a decisional step 345, a determination is made whether the incoming packet is a retransmission. If the packet is a retransmission, then in a step 351, the packet is added to the high priority (HIPRI) queue. Otherwise, in a step 353, the packet is added to the normal queue.

Fig. 3D illustrates a simplified flowchart of a timer service process useful in conjunction with data rate control process illustrated by Fig. 3C in a specific embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3D shows a step 361 invoked whenever a timer set in step 357 expires. In step 341, the byte counter is cleared and the local maximum is set to the maximum MAX added to EXTRA. Then, in a decisional step 363, the high priority (HIPRI) queue is checked to determine whether packets are on the queue. If there are packets on the HIPRI queue, then processing continues with a decisional step 365. Otherwise, processing continues with a decisional step 367.

If there are packets on the HIPRI queue, then in a decisional step 365, a determination is made whether the byte counter and the packet length are less than a local maximum. If this is indeed the case, then processing continues with a step 369 in which the packet is dequeued from the HIPRI queue and sent. Also, the byte counter is
5 incremented to reflect the length of the data in the packet. Otherwise, if decisional step 365 determines that the byte counter and the packet length exceed the local maximum, then in a step 371 the timer is restarted and processing returns to the invoking process. Processing can continue when the timer expires.

If decisional step 363 determines that there are no packets on the HIPRI
10 queue, then processing continues with a decisional step 367. Decisional step 367 determines whether there are packets on the normal queue. If there are no packets on the normal queue, then processing continues with a step 375, in which the current clock tick is stored in a global cell. Otherwise, if decisional step 367 determines that packets are on the normal queue, then processing continues with a decisional step 373, in which a
15 determination is made whether the byte counter and the packet length are less than the local maximum. If this is indeed the case, then processing continues with a step 377, in which the packet is dequeued from the normal queue and sent. Also, the byte counter is incremented to reflect the length of the data in the packet. Otherwise, if decisional step 373 determines that if the byte counter and the packet length exceed the local maximum,
20 then in a step 379 the timer is restarted and processing returns to the invoking process. Processing can continue when the timer expires.

Fig. 3E illustrates a representative system overview in a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art
25 would recognize other variations, modifications, and alternatives. A client 380 initiates a connection request to a TCP server 388. Satellite gateway 384 intercepts the connection request and establishes a second connection 385 with a second Satellite gateway 386. The second satellite gateway 386 then initiates a third connection 387 with the server 388. Once connection 387 is established and that information has been communicated back to
30 the satellite gateway 384, the first TCP connection 383 can be confirmed between the client 380 and satellite gateway 384.

In some embodiments, a data flow control process 382 useful for limiting an aggregate amount of data buffered by satellite gateway 384 and/or satellite gateway

386 for TCP connections sending data to client 380 is provided. Select embodiments can control data flow on a TCP connection when data has been buffered for a client on that connection, for example. As seen in Fig. 3F, the first TCP connection 383 can be confirmed between the client 380 and satellite gateway 384 via data flow control process 382. In select embodiments, data flow control process 382 can limit the amount of memory that a client TCP connection, such as client TCP connection 383, can send to a satellite gateway, such as satellite gateway 384, in order to control the amount of data stored on the satellite gateways 384 and satellite gateway 386 on either side of connection 385. Although depicted in terms of a separate control process 382 interposed within connection 383 between client 380 and satellite gateway 384, the present invention can also be embodied in other forms. For example, control process 382 can be co-located with first satellite gateway within client 380's machine, or can reside with first satellite gateway 384 in a separate machine, and the like.

Although depicted in terms of satellite and satellite gateways, the present invention can also be embodied in other forms of network hardware. For example, first gateway 384 and second gateway 386 can be connected by a wireless network and the like.

Fig. 3G illustrates a simplified flowchart of a connection flow control selection process in a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3G illustrates an input data flow 390. A decisional step 391 determines whether the connection on which input data flow 390 is received is a new connection. If this is so, the new connection is not marked for flow control. Otherwise, in a decisional step 392, a size of the contents of a buffer associated with the connection is checked against a threshold. If the contents of the buffer exceeds the threshold, then in a step 393, the connection is selected for the flow control limit calculation. Otherwise, the connection is not selected for flow control. The steps illustrated in Fig. 3G can provide embodiments with the capability to screen out connections that are initializing or which are transferring data intermittently. Such connections need not be selected for the flow control limit calculation because they do not consume sufficiently large system resources. New connections need not be selected for the flow control limit calculation until a threshold amount of data has been buffered for the new connection. In a presently

preferable embodiment, the threshold is reached if the connection buffers the threshold amount of data in memory. In this particular embodiment, a connection that sends less than the threshold amount, waits, and then sends more data, need not be selected for the flow control limit calculation because not all of the threshold amount of data resided in memory at one time. However, in other embodiments, an average of data in the buffer, or a buffered amount of data over a particular time period could also be used as a basis for controlling flow on the connection without departing from the scope of the present invention. In a presently preferable embodiment, the threshold limit can be 50 Kbytes. However, those of ordinary skill in the art can readily apply the methods and systems of the present invention to other threshold values.

Fig. 3H illustrates a simplified flowchart of a data flow control process useful in conjunction with the satellite gateway system illustrated by Fig. 3F in a specific embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3H illustrates an input data flow 394 for a particular connection. A decisional step 395 determines whether the connection corresponding to input data 394 is being flow controlled. If this is so, then in a step 396, a round trip time (RTT) to send a byte of data across the link and back for the connection can be determined. Round trip time can be determined by input from a user, or automatically determining the time for a sample piece of data to traverse the link, and the like. Then in a step 397, a data rate, or bandwidth, for the connection is determined. Data rate can be determined by input from a user, or automatically with a sample piece of data, and the like. In a present embodiment, the user enters the rate and round-trip time as separate inputs. However, other embodiments can allow the user to enter a pre-computed bandwidth-delay product or can determine the round-trip time by probing the link. Then, in a step 398, a bandwidth delay product can be determined by multiplying the bandwidth of the link, in bytes, by the round-trip time.

In a step 399, a flow control limit can be determined for the connection. A flow control limit can be an amount of data to buffer for the connection, or the like. In a presently preferable embodiment, the flow control limit can be an amount of memory allocated to each active data connection selected during the processing detailed in Fig. 3G. The flow control limit can be computed by multiplying the bandwidth-delay product by a scaling factor and then dividing that total by the number of active data connections.

The scaling factor is chosen to permit sufficient data to be buffered in order to keep the network link fully utilized. In a presently preferable embodiment, a scaling factor of eight is chosen. However, embodiments can use other scaling factors, or scaling functions, or the combination of multiple scaling factors without departing from the scope of the present invention. Then, in a step 401, data on the connection can be limited by buffering incoming data on the connection up to the limit determined in step 399. The system can refrain from buffering more data than there is allocated space available in physical system memory.

It is noteworthy that many embodiments according to the present invention can adapt to change in the number of active connections. If one connection is active, that connection will be permitted to use the entire bandwidth. When a second connection becomes active, flow control limits can be adjusted for both the first and second connections by calculating a flow control limit whenever new data arrives. In this manner, bandwidth can be allocated among connections irrespective of the order that connections are established. Additionally, embodiments according to the present invention permit sharing of a satellite link between connections by reducing the likelihood that a burst of data from one particular connection will be queued for transmission over the link in front of data for other connections. This makes network throughput more consistent across multiple connections.

Fig. 3I illustrates a simplified structure component diagram of a representative buffer architecture of a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3I illustrates a satellite gateway 384. Satellite gateway 384 comprises a TCP buffer 405, that interfaces with a client, such as client 380 of Fig. 3E via connection 383. Further, satellite gateway 384 comprises a satellite protocol buffer 406, that interfaces with a satellite, such as satellite 101 of Fig. 3E via connection 385. TCP over satellite systems can temporarily store data in one or more buffers controlled by the TCP protocol layer and the satellite protocol layer in satellite gateway 384. Data to an endpoint, such as client 380 of Fig. 3E, for example, can be delivered from one or more TCP buffers, such as TCP buffer 405. As space becomes available in TCP buffer 405, data can be transferred from satellite protocol buffer 406 to TCP buffer 405. As space becomes available in satellite protocol buffer 406, the satellite

protocol can permit more data to be sent across the satellite link, such as satellite link 385 in Fig. 3E, for example. In a representative example of an application of an embodiment according to the present invention, a satellite gateway, such as satellite gateway 384 can feed data to a plurality of modems. The modem links cannot accept data at the same rate as the satellite link. In such a situation, techniques according to the present invention can circumvent expending relatively large amounts of buffer space for holding data for the modems.

Fig. 3J illustrates a simplified flowchart of a connection memory management process in a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3J illustrates output data 402, that can reside in satellite protocol buffer 406 of Fig. 3I, for example. When data is not being delivered from the TCP layer to the end client at an acceptable rate, congestion in the satellite link can occur. This congestion can be detected by the satellite protocol in a decisional step 403, wherein the satellite protocol determines whether a state of the satellite protocol buffer 406 has become saturated. In some embodiments, a threshold can be established for the buffer. Exceeding this threshold can indicate a saturation condition is present. If the protocol buffer has become saturated, then in a step 404, the satellite protocol reduces a maximum receive window for data that will be accepted across the satellite link 385 by a reduction factor. In a presently preferable embodiment, a reduction factor of 50% can be used. In a particular embodiment, an minimum value for the receive window can be configured to assure that the window is not reduced to zero. Otherwise, if a client is receiving at a data rate comparable to the data rate of the system, then the saturation condition will not be reached.

In a particular embodiment according to the present invention, a connection to a receiver, such as client 384, for example, that has a relatively slow data rate, can start with a full-sized receive window. Using the techniques according to the present invention, the receive window can be resized to achieve a match between the data rate of the sender and the data rate of the receiver. Thus, in some embodiments, new connections can have the capability of consuming a full window of memory before the system detects the lower data rate condition and restricts how much additional new data will be buffered for the connection. In a presently preferable embodiment, the techniques

of the present invention can be applied to each satellite connection separately, so that one slow receiver will not degrade performance of other connections over the same link.

Fig. 7 illustrates a representative apparatus in a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 7 depicts a block diagram of a protocol gateway in a particular embodiment according to the present invention, such as protocol gateway 111A, or the like. Protocol gateway 111A includes a bus 112 which interconnects major subsystems such as a central processor 114, a system memory 116 (typically RAM), and a number of optional external devices such as a display screen 124 via a display adapter 126, a keyboard 132 and a mouse 146 via an I/O controller 118, a floppy disk drive 136 operative to receive a floppy disk 138. Storage Interface 134 may act as a storage interface to a fixed disk drive 144. An optional CD-ROM player 140 operative to receive a CD-ROM 142 can also be included. Other forms of storage, such as a flash memory 147 can also be included. A network interface 148A can provide a direct connection to a remote server via a telephone link or to the Internet using a common protocol such as TCP/IP or the like. Network interface 148A can interface to one or more outside networks, which can employ any network format known in the art, such as Ethernet, Token Ring, ATM, IEEE 802.3, X.25, Serial Link Internet Protocol (SLIP), FDDI and the like. Network interface 148B may also connect to a satellite gateway, such as satellite gateway 109A, or other network interconnecting devices or computer systems. Network interface 148A can interface to one or more outside networks, which can employ any network format known in the art, such as Ethernet, Token Ring, ATM, IEEE 802.3, X.25, Serial Link Internet Protocol (SLIP), FDDI and the like. Many other devices or subsystems (not shown) may be connected in a similar manner. Also, it is not necessary for all of the devices shown in Fig. 7 to be present to practice the present invention, as discussed below. The devices and subsystems may be interconnected in different ways from that shown in Fig. 7 in various embodiments. Code to implement the present invention, may be operably disposed or stored in computer-readable storage media such as system memory 116, fixed disk 144, CD-ROM 140, or floppy disk 138.

System 111A is merely one example of a configuration that embodies the present invention. It will be readily apparent to one of ordinary skill in the art that many

system types, configurations, and combinations of the above devices are suitable for use in light of the present disclosure.

Although the above has generally described the present invention according to specific systems, methods and embodiments, the present invention has a much broader range of applicability. In particular, the present invention is not limited to satellite telecommunications, but can be applied to any networking situation where an improved or optimized protocol is desired for use over a specific portion of the network, and the end systems cannot be updated to use the improved protocol. Thus, in some embodiments, satellite gateways could provide access to wireless or cabled networks and internetworks of all kinds. Of course, one of ordinary skill in the art would recognize other variations, modifications, and alternatives.

Experiment:

To prove the principles and operation of the present system, experiments have been performed. In the present experiments, we ran a series of single-client FTP file transfer throughput tests to compare performance of conventional TCP and the present invention over a wide variety of different TCP windows sizes, link bandwidths, round trip delay times, and bit error rates. In one experiment, "Window Size" was compared against "Throughput," as shown by the simplified diagram 400 of Fig. 4, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Without performance enhancements or TCP window size tuning, most clients were limited to a throughput rate of less than 100 Kbps when the round trip time was 540 ms. As the test data in Fig. 4 show, even clients with a 32 KB window were able to reach a throughput of a mere 400 Kbps. The present system allowed the client to take advantage of the available bandwidth regardless of the window size of the client or server.

In an alternative experiment, a "Round Trip Delay" was analyzed against "Throughput," as shown in a simplified diagram 500 of Fig. 5, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Here, the present invention removed the dependency of TCP on the round trip time of the network. Performance was measured over symmetric 2 Mbps and 10 Mbps links with no errors. These results illustrated that TCP is able to fully saturate terrestrial low-delay

networks, but as the delay increases, TCP performance dropped rapidly. In contrast, a network which used the present system was able to maintain near complete usage of the link regardless of the round trip time of the network.

In an alternative experiment, a "Bit Error Rate" was monitored against "Throughput," which is shown by the simplified diagram 600 of Fig. 6, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Networks which incorporated the present invention were also sensitive to the bit error rate of the link. The diagram shows throughput as a function of the bit error rate for a symmetric 10 Mbps satellite network with a 540 ms round trip time and a TCP window size of 1 MB. Even at low error rates, TCP was able to deliver a mere 1.2 Mbps. At an error rate of 1.0×10^{-5} , TCP's throughput dropped to 0.05 Mbps. A network which used the present system was able to fully saturate the link at low error rates and delivered 2.7 Mbps even at an error rate of 1.0×10^{-5} .

While the above is a full description of the specific embodiments, various modifications, alternative constructions and equivalents may be used. For example, the above has generally been described in terms of using XTP as a satellite protocol. Other types of protocols may be available. For example, the protocol can include a modified TCP or the like. Therefore, the above description and illustrations should not be taken as limiting the scope of the present invention which is defined by the appended claims.

WHAT IS CLAIMED IS:

1 1. A communication apparatus for transmitting packetized
2 information, said information comprising a plurality of packets, each of said packets
3 comprising data and a header, over a satellite link in a telecommunications system, said
4 system comprising a client, selected from a plurality of potential clients, a server, selected
5 from a plurality of potential servers, a first gateway, connected to said client by a first
6 telecommunications link, a second gateway, connected to said server by a second
7 telecommunications link, a third telecommunications link connecting said first gateway to
8 said second gateway, said apparatus comprising:

9 a network interface for linking said first gateway with said client;
10 a satellite gateway interface;
11 a system memory; and
12 a bus interconnecting said network interface, said satellite gateway
13 interface, and said system memory with a processor, said processor operatively disposed
14 to:

15 intercept a connection with said server, said communication initiated by
16 said client;

17 establish a connection between said first gateway and said second gateway
18 over said telecommunications link;

19 provide a bi-directional flow of information from said client to said server
20 and from said server to said client using said connection between said first gateway and
21 said second gateway, wherein said providing a bi-directional flow occurs transparently to
22 said client and said server.

1 2. The apparatus of claim 1 wherein said processor is further
2 operatively disposed to convert said information at said first gateway from a first protocol
3 into a second protocol for transmission over said telecommunications link; and

4 convert said second protocol into said first protocol at said second
5 gateway.

1 3. The apparatus of claim 2 wherein the first protocol comprises TCP
2 and said second protocol comprises XTP.

1 4. The apparatus of claim 2 wherein said second protocol is more
2 suitable for transmission over a satellite link than using a TCP protocol.

1 5. A communication apparatus comprising:
2 a TCP interface;
3 a satellite gateway interface;
4 a system memory;
5 a bus interconnecting said TCP interface, said satellite gateway interface
6 and said system memory with a processor, said processor operatively disposed to:
7 intercept a first communication connection between a client and a server;
8 form a second communication connection between a first satellite gateway
9 and a second satellite gateway that is over a satellite link;
10 transmit information describing said first connection to said second
11 satellite gateway; and
12 form a third communication connection between said second satellite
13 gateway and a destination server using said information describing said first connection
14 wherein said forming said second connection and forming said third connection occur
15 transparently to said client and said server.

1 6. The apparatus of claim 5 wherein said information comprises a
2 client address and a destination server address.

1 7. The apparatus of claim 5 wherein said processor is further
2 operatively disposed to transmit a response from said second satellite gateway to said first
3 satellite gateway when said third communication connection with said destination server
4 occurs.

1 8. The apparatus of claim 5 wherein said processor is further
2 operatively disposed to transmit a response from said first satellite gateway to said client
3 when said third communication connection with said destination server occurs.

1 9. The apparatus of claim 5 wherein said processor is further
2 operatively disposed to transmit a failure response from said first satellite gateway to said
3 client when said third communication connection is lost.

1 10. A apparatus for establishing a communication between a first
2 computer and one of a plurality of second computers, said apparatus comprising:
3 a network interface;
4 a satellite gateway interface;
5 a system memory;
6 a bus interconnecting said network interface, said satellite gateway
7 interface and said system memory with a processor, said processor operatively disposed
8 to:

9 provide a first protocol and a second protocol, wherein at least one of said
10 plurality of second computers is able to communicate with said first computer using said
11 first protocol and at least one of said plurality of second computers is able to
12 communicate with said first computer using said second protocol;

13 determine whether a connection can be established between said first
14 computer and at least one of said plurality of second computers using said first protocol;

15 if said connection cannot be established, then establish a connection
16 between said first computer and said at least one of said plurality of second computers
17 using said second protocol.

1 11. The apparatus of claim 10 wherein each of said plurality of second
2 computers is able to communicate with said first computer using said second protocol.

1 12. The apparatus of claim 10 wherein said first protocol is XTP.

1 13. The apparatus of claim 10 wherein said second protocol is TCP/IP.

1 14. The apparatus of claim 10 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 7.5 times greater than said second throughput at a bit error rate of 1×10^{-7}

1 15. The apparatus of claim 10 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 10 times greater than said second throughput at a bit error rate of 1×10^{-6}

1 16. The apparatus of claim 10 wherein said first protocol has a
2 throughput of at least 95% of an available bandwidth at a bit error rate of 1×10^{-8} .

1 17. In a communications system for providing transport of packetized
2 information in a TCP protocol between a client and a server, a method of converting said
3 system to a system for providing transport of packetized information in a satellite
4 protocol, said method comprising:

5 installing a first gateway, said first gateway operatively disposed to
6 intercept connections from said client to said server; said first gateway further operatively
7 disposed to convert said packetized information from TCP protocol to satellite protocol;

8 installing a second gateway, said second gateway disposed to establish a
9 connection with said server, said second gateway further operatively disposed to convert
10 said packetized information from satellite protocol to TCP protocol, said first gateway
11 and said second gateway operatively disposed to establish a connection over a common
12 telecommunications link.

1 18. A communication system for transmitting packetized information,
2 said information comprising a plurality of packets, each of said packets comprising data
3 and a header, over a satellite link, said system comprising:

4 a client, selected from a plurality of potential clients;

5 a server, selected from a plurality of potential servers;

6 a first gateway, connected to said client by a first telecommunications link;

7 a second gateway, connected to said server by a second
8 telecommunications link;

9 a third telecommunications link connecting said first gateway to said
10 second gateway;

11 code for intercepting a connection attempt with said server, said
12 connection attempt initiated by said client;

13 code for establishing a connection between said first gateway and said
14 second gateway over said third telecommunications link;

15 code for providing a bi-directional flow of information from said client to
16 said server and from said server to said client using said connection between said first
17 gateway and said second gateway, wherein said providing a bi-directional flow occurs
18 transparently to said client and said server; and

19 a computer readable storage medium for storing said codes.

1 19. The system of claim 18 further comprising:

code for converting said information at said first gateway from a first protocol into a second protocol for transmission over said telecommunications link; and code for converting said second protocol into said first protocol at said second gateway.

20. The system of claim 19 wherein said code for converting comprises removing said header to leave said data substantially intact.

21. The system of claim 19 wherein said code for converting comprises code for removing said header to leave said data substantially intact and code for encapsulating said data using a satellite protocol header.

22. The system of claim 21 wherein said data is a portion of said flow of information.

23. The system of claim 18 further comprising code for receiving said flow of information by said second gateway over said telecommunications link.

24. A communication system comprising:
code for intercepting a first communication connection between a client and a server;

code for forming a second communication connection between a first satellite gateway and a second satellite gateway that is over a satellite link;

code for transmitting information describing said first connection to said second satellite gateway;

code for forming a third communication connection between said second satellite gateway and a destination server using said information describing said first connection wherein said code for forming said second connection and code for forming said third connection execute transparently to said client and said server; and

a computer readable storage medium for holding said codes.

25. The system of claim 24 wherein said information comprises a client address and a destination server address.

26. A system for establishing a communication between a first computer and one of a plurality of second computers, said system comprising:

code for providing a first protocol and a second protocol, wherein at least one of said plurality of second computers is able to communicate with said first computer using said first protocol and at least one of said plurality of second computers is able to communicate with said first computer using said second protocol;

code for determining whether a connection can be established between said first computer and at least one of said plurality of second computers using said first protocol;

code for establishing a connection between said first host computer and said at least one of said plurality of second host computers using said second protocol, if said code for determining determines that a connection cannot be established between said first computer and said second computer using said first protocol; and

a computer readable storage medium for holding said codes.

27. The system of claim 26 wherein each of said plurality of second computers is able to communicate with said first computer using said second protocol.

28. The system of claim 26 wherein said first protocol is XTP.

29. The system of claim 26 wherein said second protocol is TCP/IP.

30. A communication method for transmitting information, said information comprising a plurality of packets, each of said packets comprising data and a header, in a system comprising a client, selected from a plurality of potential clients;

a server, selected from a plurality of potential servers;

a first gateway, connected to said client by a first telecommunications link;

a second gateway, connected to said server by a second telecommunications link;

a third telecommunications link connecting said first gateway to said second gateway, said method comprising:

intercepting a connection attempt with said server, said connection attempt initiated by said client;

establishing a connection between said first gateway and said second gateway over said third telecommunications link;

providing a bi-directional flow of information from said client to said server and from said server to said client using said connection between said first gateway

16 and said second gateway, wherein said providing a bi-directional flow occurs
17 transparently to said client and said server.

1 31. The method of claim 30 further comprising:
2 converting said information at said first gateway from a first protocol into
3 a second protocol for transmission over said third telecommunications link; and
4 converting said second protocol into said first protocol at said second
5 gateway.

1 32. The method of claim 31 wherein the first protocol comprises TCP
2 and the second protocol comprises XTP.

1 33. The method of claim 31 wherein said second protocol is more
2 suitable for transmission over a satellite link than using a TCP protocol.

1 34. The method of claim 31 wherein said converting comprises
2 removing said header to leave said data substantially intact.

1 35. The method of claim 31 wherein said converting comprises
2 removing said header to leave said data substantially intact and encapsulating said data
3 using a satellite protocol header.

1 36. The method of claim 35 wherein said data is a portion of said flow
2 of information.

1 37. The method of claim 30 further comprising receiving said flow of
2 information by a gateway over said third telecommunications link.

1 38. A communication method comprising:
2 intercepting a first communication connection between a client and a
3 server;
4 forming a second communication connection between a first satellite
5 gateway and a second satellite gateway that is over a satellite link;
6 transmitting information describing said first connection to said second
7 satellite gateway; and
8 forming a third communication connection between said second satellite
9 gateway and a destination server using said information describing said first connection

10 wherein said forming said second connection and forming said third connection occurs
11 transparently to said client and said server.

1 39. The method of claim 38 wherein said information comprises a
2 client address and a destination server address.

1 40. The method of claim 38 further comprising transmitting a response
2 from said second satellite gateway to said first satellite gateway when said third
3 communication connection with said destination server occurs.

1 41. The method of claim 38 further comprising transmitting a response
2 from said first satellite gateway to said client when said third communication connection
3 with said destination server occurs.

1 42. The method of claim 38 further comprising transmitting a failure
2 response from said first satellite gateway to said client when said third communication
3 connection is lost.

1 43. A method for establishing a communication between a first
2 computer and one of a plurality of second computers, said method comprising:
3 providing a first protocol and a second protocol, wherein at least one of
4 said plurality of second computers is able to communicate with said first computer using
5 said first protocol and at least one of said plurality of second computers is able to
6 communicate with said first computer using said second protocol;
7 determining whether a connection can be established between said first
8 computer and at least one of said plurality of second computers using said first protocol;
9 establishing a connection between said first computer and said at least one
10 of said plurality of second computers using said second protocol if said determining step
11 determines that a connection cannot be established between said first computer and said
12 second computer using said first protocol.

1 44. The method of claim 43 wherein each of said plurality of second
2 computers is able to communicate with said first computer using said second protocol.

1 45. The method of claim 43 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 7.5 times greater than said second throughput at a bit error rate of 1×10^{-7}

1 46. The method of claim 43 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 10 times greater than said second throughput at a bit error rate of 1×10^{-6} .

1 47. The method of claim 43 wherein said first protocol has a
2 throughput of at least 95% of an available bandwidth at a bit error rate of 1×10^{-8} .

1 48. A method for managing memory for buffering information, said
2 information comprising a plurality of packets, each of said packets comprising data and a
3 header, in a system comprising a client, selected from a plurality of potential clients;

4 a server, selected from a plurality of potential servers;

5 a first gateway, connected to said client by a first telecommunications link,
6 said first gateway comprising a first buffer and a second buffer;

7 a second gateway, connected to said server by a second
8 telecommunications link;

9 a third telecommunications link connecting said first gateway to said
10 second gateway, said method comprising:

11 intercepting a connection attempt with said server, said connection attempt
12 initiated by said client, said client capable of establishing connections at a first
13 characteristic data rate;

14 establishing a connection between said first gateway and said second
15 gateway over said third telecommunications link, said connection having a second
16 characteristic data rate; wherein said first buffer stores information received over said
17 third telecommunications link and said second buffer stores information for said client;

18 determining a state in said first buffer, said state indicating that said
19 second characteristic data rate in said third telecommunications link is greater than said
20 first characteristic data rate for said client; and

21 setting a receive window size for said client at said first gateway to reduce
22 an amount of data received over said third telecommunications link for said client.

1 49. The method of claim 48 wherein said state comprises a buffer full
2 condition.

1 50. The method of claim 48 wherein said state comprises a threshold
2 value for buffer contents has been reached.

1 51. The method of claim 48 wherein said setting said receive window
2 size comprises reducing a receive window size by about 50%.

1 52. The method of claim 48 wherein said connection attempt by said
2 client with said server comprises a TCP protocol.

1 53. The method of claim 48, wherein said connection between said
2 first gateway and said second gateway comprises a satellite protocol.

1 54. The method of claim 53 wherein said satellite protocol further
2 comprises an XTP protocol.

1 55. A computer program product for managing memory for buffering
2 information, said information comprising a plurality of packets, each of said packets
3 comprising data and a header, in a system comprising a client, selected from a plurality of
4 potential clients;

5 a server, selected from a plurality of potential servers;

6 a first gateway, connected to said client by a first telecommunications link,
7 said first gateway comprising a first buffer and a second buffer;

8 a second gateway, connected to said server by a second
9 telecommunications link;

10 a third telecommunications link connecting said first gateway to said
11 second gateway, said computer program product comprising:

12 code for intercepting a connection attempt with said server, said
13 connection attempt initiated by said client, said client capable of establishing connections
14 at a first characteristic data rate;

15 code for establishing a connection between said first gateway and said
16 second gateway over said third telecommunications link, said connection having a second

characteristic data rate; wherein said first buffer stores information received over said third telecommunications link and said second buffer stores information for said client; code for determining a state in said first buffer, said state indicating that said second characteristic data rate in said third telecommunications link is greater than said first characteristic data rate for said client; code for setting a receive window size for said client at said first gateway to reduce an amount of data received over said third telecommunications link for said client; and a computer readable storage medium for holding the codes.

56. The computer program product of claim 55 wherein said state comprises a buffer full condition.

57. The computer program product of claim 55 wherein said state comprises a threshold value for buffer contents has been reached.

58. The computer program product of claim 55 wherein said code for setting said receive window size comprises code for reducing a receive window size by about 50%.

59. The computer program product of claim 55 wherein said connection attempt by said client with said server comprises a TCP protocol.

60. The computer program product of claim 55, wherein said connection between said first gateway and said second gateway comprises a satellite protocol.

61. The computer program product of claim 60 wherein said satellite protocol further comprises an XTP protocol.

62. A system for managing memory for buffering a flow of information over a satellite, said system comprising:
a client, selected from a plurality of potential clients;
a server, selected from a plurality of potential servers;
a first gateway, connected to said client by a first telecommunications link, said first gateway comprising a first buffer and a second buffer;

7 a second gateway, connected to said server by a second
8 telecommunications link;

9 a third telecommunications link connecting said first gateway to said
10 second gateway; wherein said first gateway is operatively disposed to:

11 intercept a connection attempt with said server, said connection attempt
12 initiated by said client, said client capable of establishing connections at a first
13 characteristic data rate;

14 establish a connection between said first gateway and said second gateway
15 over said third telecommunications link, said connection having a second characteristic
16 data rate; wherein said first buffer stores information received over said third
17 telecommunications link and said second buffer stores information for said client;

18 determine a state in said first buffer, said state indicating that said second
19 characteristic data rate in said third telecommunications link is greater than said first
20 characteristic data rate for said client; and

21 set a receive window size for said client at said first gateway to reduce an
22 amount of data received over said third telecommunications link for said client.

1 63. The system of claim 62 wherein said state comprises a buffer full
2 condition.

1 64. The system of claim 62 wherein said state comprises a threshold
2 value for buffer contents has been reached.

1 65. The system of claim 62 wherein said setting said receive window
2 size comprises reducing a receive window size by about 50%.

1 66. The system of claim 62 wherein said connection attempt by said
2 client with said server comprises a TCP protocol.

1 67. The system of claim 62 wherein said connection between said first
2 gateway and said second gateway comprises a satellite protocol.

1 68. The system of claim 67 wherein said satellite protocol further
2 comprises an XTP protocol.

1 69. A method for controlling the flow of information, said information
2 comprising a plurality of packets, each of said packets comprising data and a header, in a
3 system comprising a client, selected from a plurality of potential clients;
4 a server, selected from a plurality of potential servers;
5 a first gateway, connected to said client by a first telecommunications link;
6 a second gateway, connected to said server by a second
7 telecommunications link;
8 a third telecommunications link connecting said first gateway to said
9 second gateway, said method comprising:
10 intercepting a connection attempt with said server, said connection attempt
11 initiated by said client;
12 establishing a connection between said first gateway and said second
13 gateway over said third telecommunications link;
14 determining a round trip time for data to travel from said first gateway to
15 said second gateway over said connection;
16 determining a data rate for said connection;
17 determining a bandwidth delay product from said round trip time and said
18 data rate;
19 determining a flow control limit for said connection from said bandwidth
20 delay product; and
21 limiting data transfer over said connection based upon said flow control
22 limit.

1 70. The method of claim 69 wherein said connection uses a satellite
2 protocol.

1 71. The method of claim 69 wherein said determining a round trip time
2 for said connection further comprises obtaining said round trip time from a user.

1 72. The method of claim 69 wherein said determining a data rate for
2 said connection further comprises obtaining said round trip time from a user.

1 73. The method of claim 69 wherein said determining a bandwidth
2 delay product for said connection further comprises multiplying said round trip time and
3 said data rate.

1 74. The method of claim 69 wherein said determining a flow control
2 limit for said connection further comprises multiplying said bandwidth delay product by a
3 scaling factor and dividing by a number of connections for said link.

1 75. The method of claim 74 wherein said scaling factor is eight.

1 76. The method of claim 69 further comprising:
2 determining from at least one of a plurality of connections at said first
3 gateway ones that can benefit from flow control.

1 77. The method of claim 76 wherein said determining comprises
2 selecting from said plurality of connections ones that are not initializing.

1 78. The method of claim 76 wherein said determining comprises
2 selecting from said plurality of connections ones that are buffering data above a threshold.

1 79. A system for controlling the flow of information over a satellite,
2 said system comprising:

3 a client, selected from a plurality of potential clients;
4 a server, selected from a plurality of potential servers;
5 a first gateway, connected to said client by a first telecommunications link;
6 a second gateway, connected to said server by a second

7 telecommunications link;

8 a third telecommunications link connecting said first gateway to said
9 second gateway; wherein said client is operatively disposed to

10 determine a round trip time for data to traverse from said first gateway
11 over said third telecommunications link to said second gateway over a connection;

12 determine a data rate for said connection;

13 determine a bandwidth delay product from said round trip time and said
14 data rate;

15 determine a flow control limit for said connection from said bandwidth
16 delay product; and

17 limit data transfer over said connection based upon said flow control limit.

1 80. The system of claim 79 wherein said connection uses a satellite
2 protocol.

1 81. The system of claim 79 wherein said determine a round trip time
2 for said connection further comprises obtaining said round trip time from a user.

1 82. The system of claim 79 wherein said determine a data rate for said
2 connection further comprises obtaining said round trip time from a user.

1 83. The system of claim 79 wherein said determine a bandwidth delay
2 product for said connection further comprises multiplying said round trip time and said
3 data rate.

1 84. The system of claim 79 wherein said determine a flow control limit
2 for said connection further comprises multiplying said bandwidth delay product by a
3 scaling factor and dividing by a number of connections for said link.

1 85. The system of claim 84 wherein said scaling factor is eight.

1 86. The system of claim 84 wherein said client is further operatively
2 disposed to:
3 determine from at least one of a plurality of connections ones that can
4 benefit from flow control.

1 87. The system of claim 86 wherein said determine comprises selecting
2 from said plurality of connections ones that are not initializing.

1 88. The system of claim 86 wherein said determine comprises selecting
2 from said plurality of connections ones that are buffering data above a threshold.

1 89. A computer program product for controlling the flow of
2 information over a data link, said data link having at least one of a plurality of
3 connections, said computer program product comprising:

4 code for determining a round trip time for data to traverse said data link on
5 said connection;

6 code for determining a data rate for said connection;

code for determining a bandwidth delay product from said round trip time and said data rate;

code for determining a flow control limit for said connection from said bandwidth delay product;

code for limiting data transfer over said connection based upon said flow control limit; and

a computer readable storage medium for containing the codes.

90. The computer program product of claim 89 wherein said connection comprises a TCP connection.

91. The computer program product of claim 89 wherein said code for determining a round trip time for said connection further comprises code for obtaining said round trip time from a user.

92. The computer program product of claim 89 wherein said code for determining a data rate for said connection further comprises code for obtaining said round trip time from a user.

93. The computer program product of claim 89 wherein said code for determining a bandwidth delay product for said connection further comprises code for multiplying said round trip time and said data rate.

94. The computer program product of claim 89 wherein said code for determining a flow control limit for said connection further comprises code for multiplying said bandwidth delay product by a scaling factor and dividing by a number of connections for said link.

95. The computer program product of claim 94 wherein said scaling factor is eight.

96. The computer program product of claim 89 further comprising:
code for determining from at least one of a plurality of connections ones that can benefit from flow control.

1 97. The computer program product of claim 96 wherein said code for
2 determining comprises code for selecting from said plurality of connections ones that are
3 not initializing.

1 98. The computer program product of claim 96 wherein said code for
2 determining comprises code for selecting from said plurality of connections ones that are
3 buffering data above a threshold.

1 99. A method for controlling the flow of information over a satellite
2 data link, said satellite data link having at least one of a plurality of connections, said
3 method comprising:


4 determining a round trip time for data to travel from a first node on said
5 satellite data link to a second node on said satellite data link over said connection;

6 determining a data rate for said connection;

7 determining a bandwidth delay product from said round trip time and said
8 data rate;

9 determining a flow control limit for said connection from said bandwidth
10 delay product; and

11 limiting data transfer over said connection based upon said flow control
12 limit.

**FIG. 1**




FIG. 2


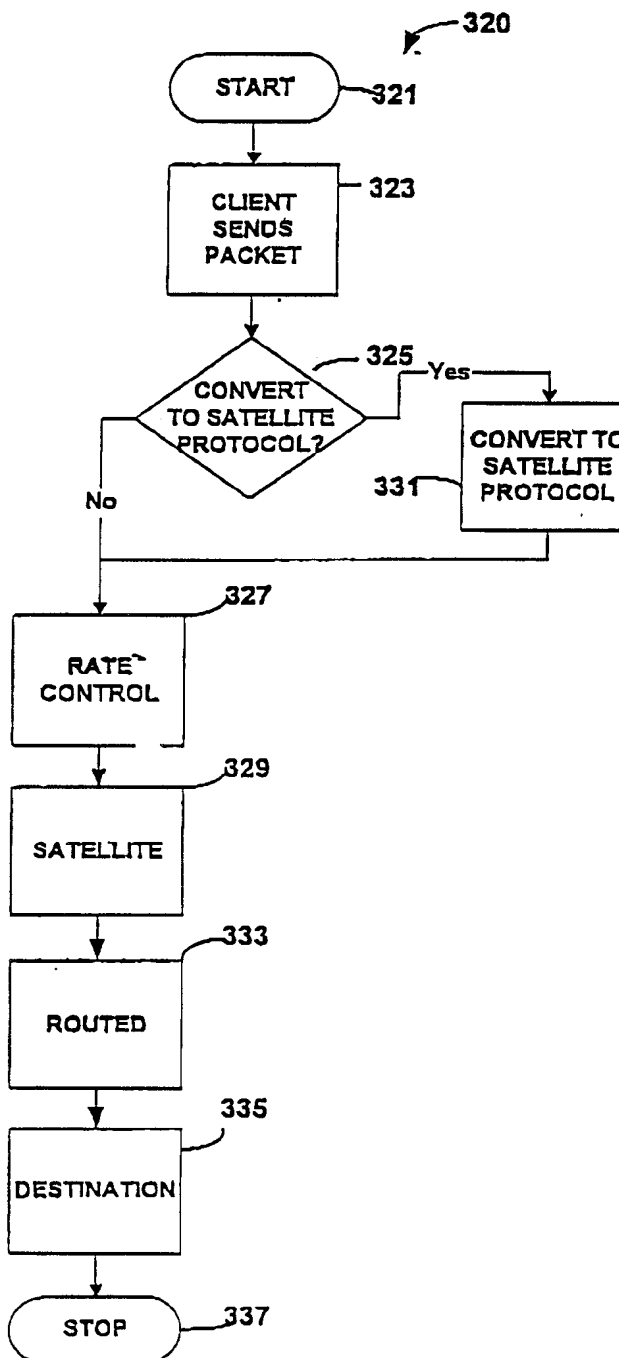



FIG. 2A



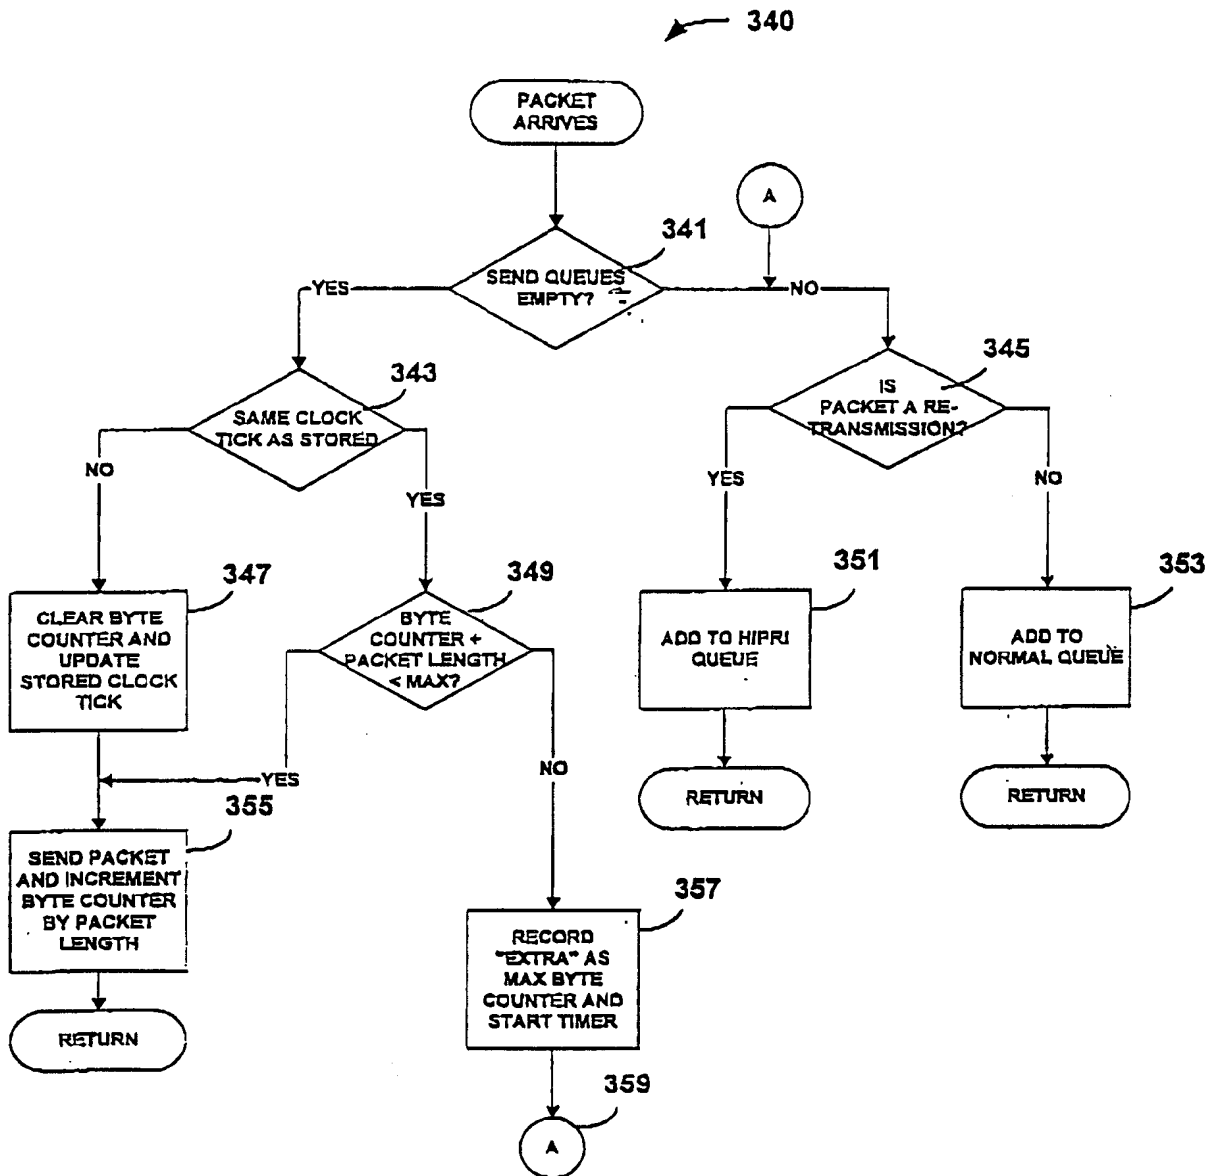


Fig. 3C

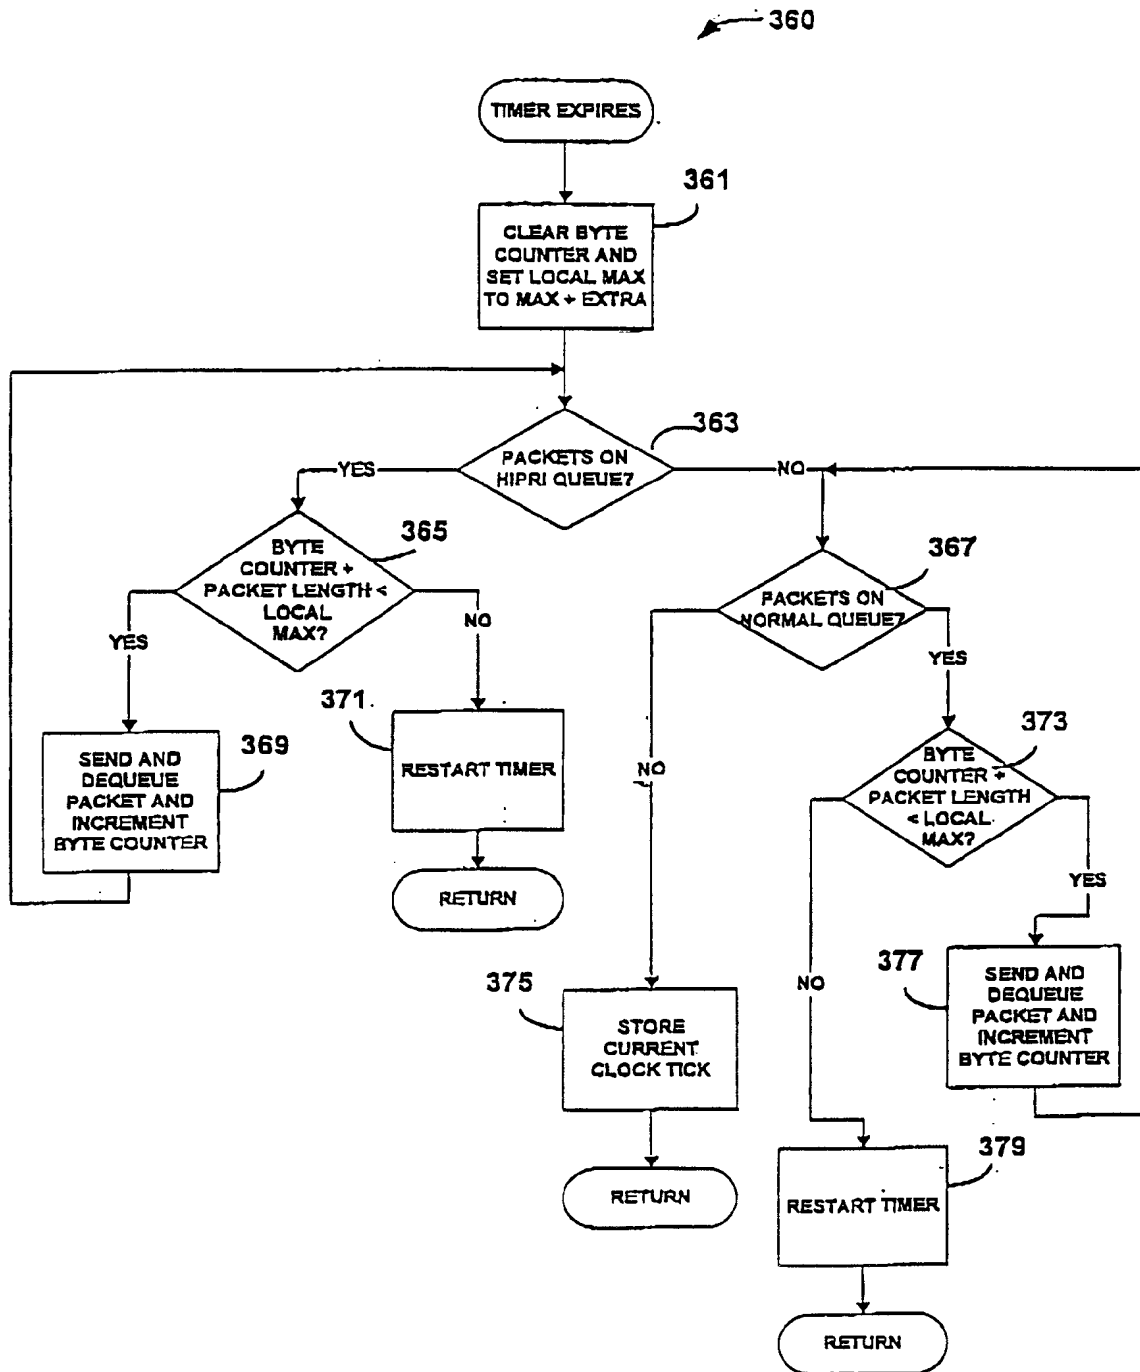


Fig. 3D




Fig. 3E




Fig. 3F

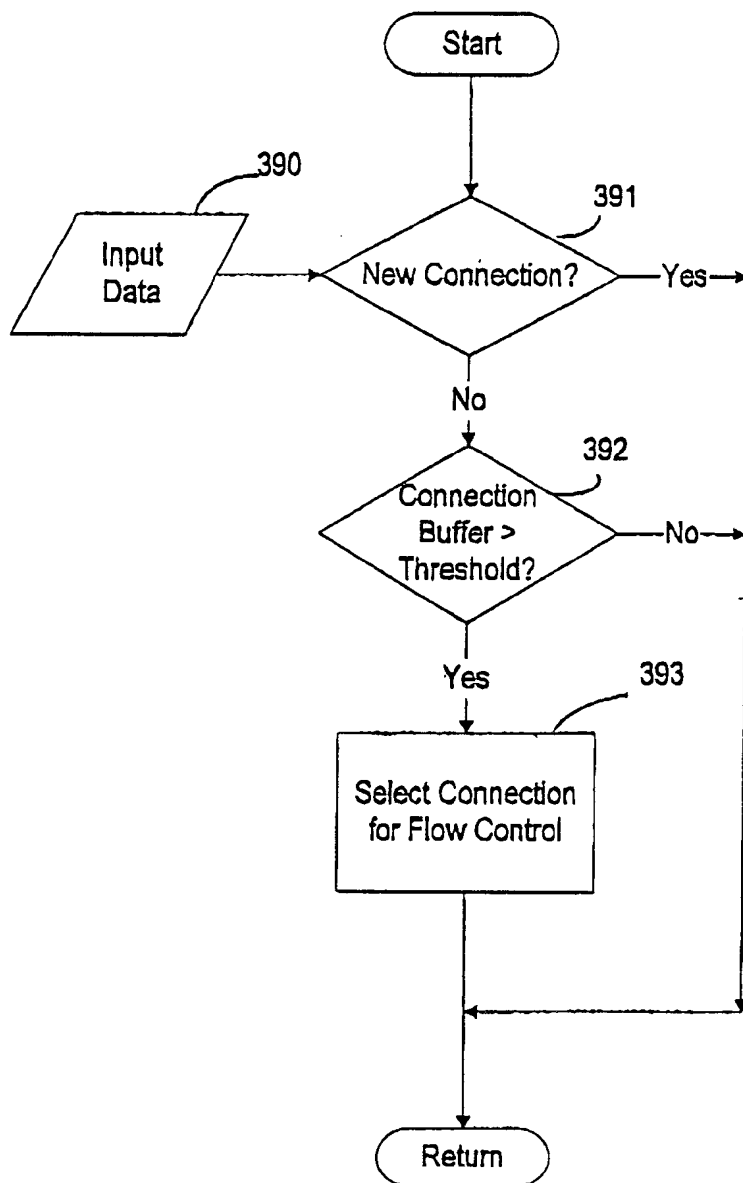


FIG. 3G




FIG. 3H




FIG. 3I

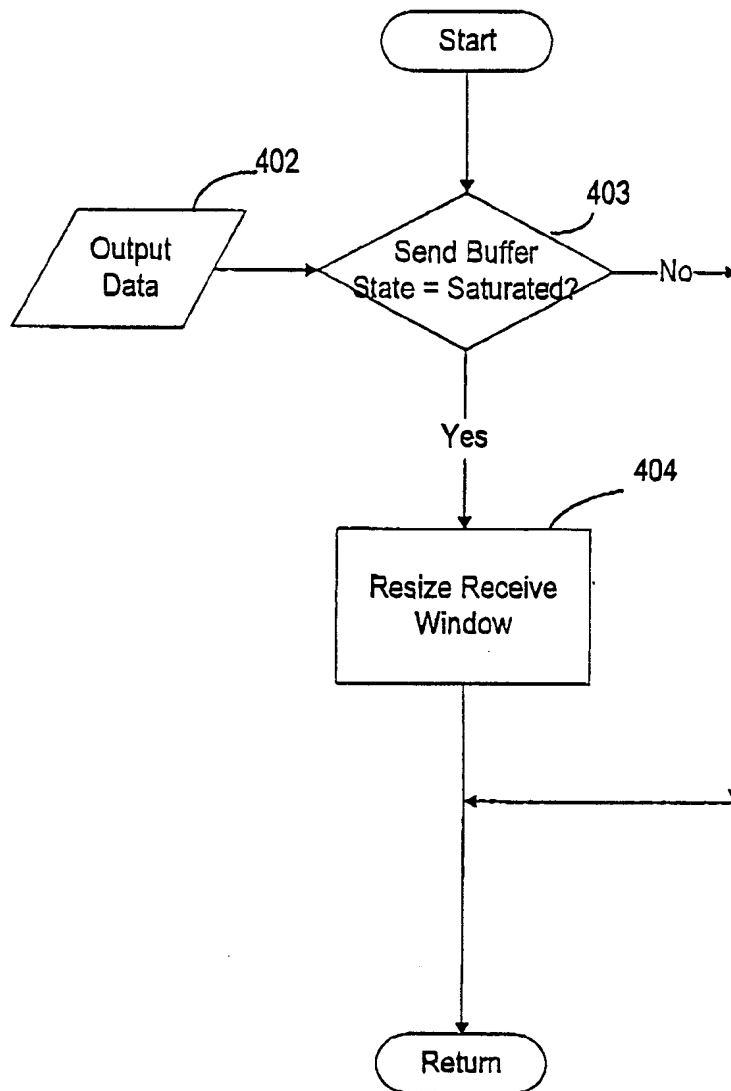




FIG. 3J




Window Size vs. Throughput

FIG. 4



Delay vs. Throughput

FIG. 5



Bit Error Rate vs. Throughput (10Mbps Link)

FIG. 6

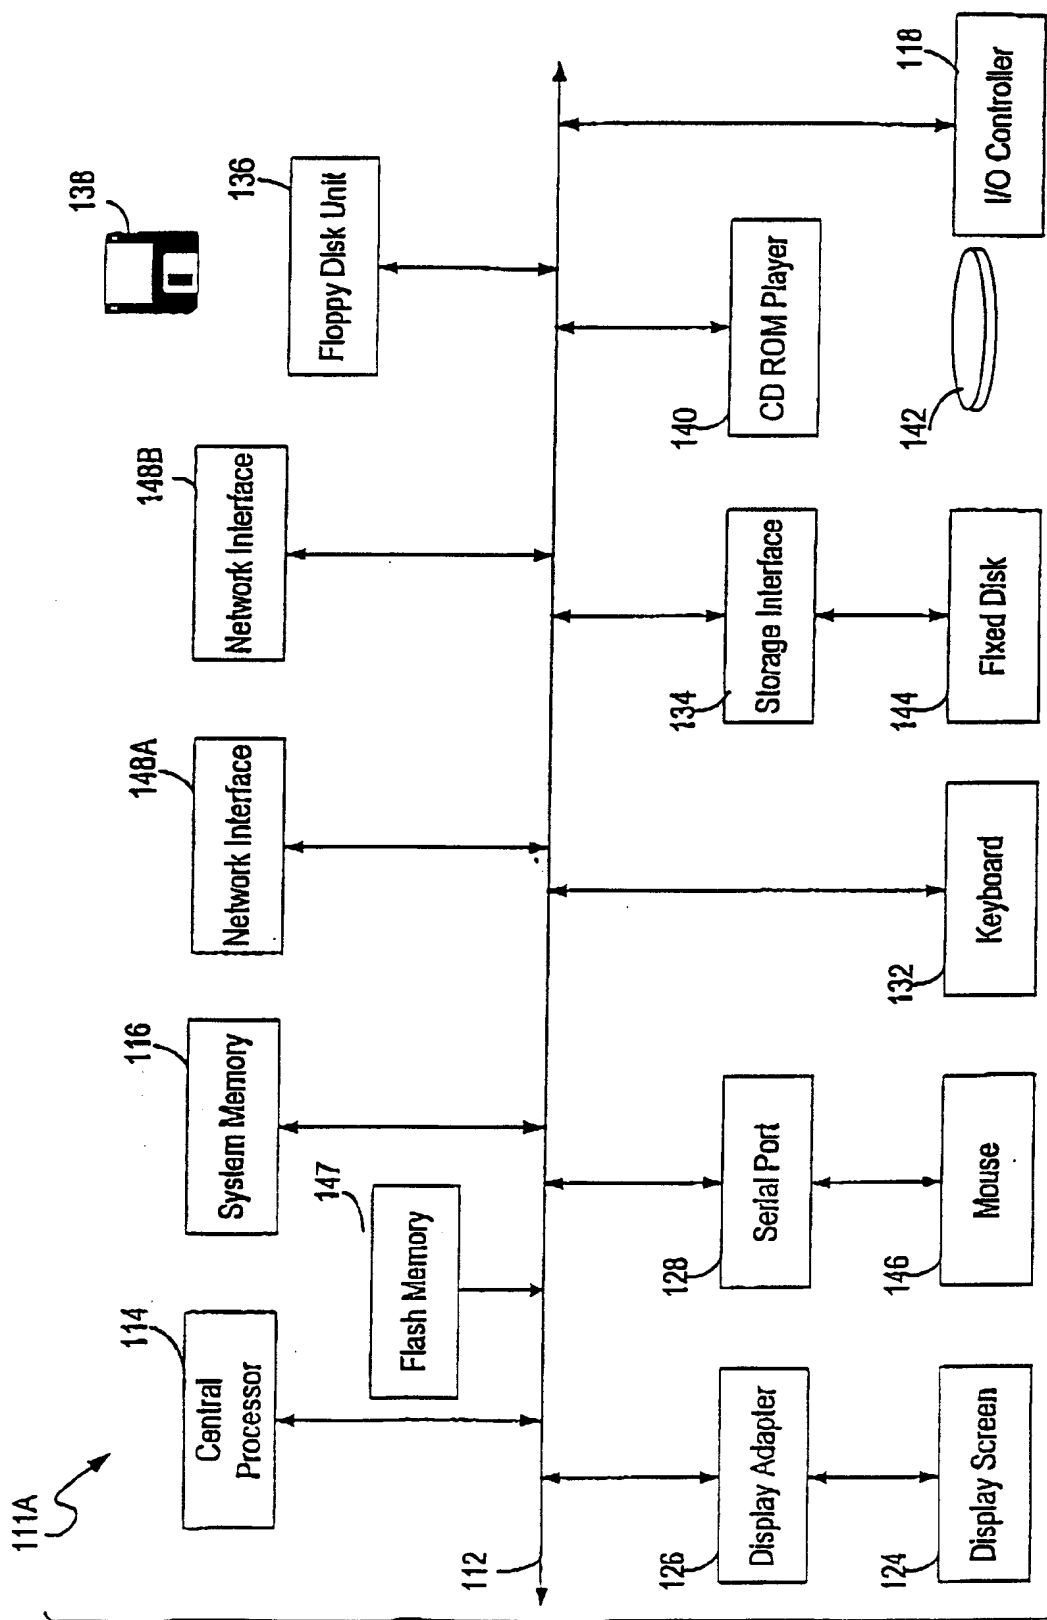


Fig. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/02891

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :IPC(7): G06F 9/44, 13/14; H04L 12/26, 46, 66; H04Q 7/00

US CL :Please See Extra Sheet.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : Please See Extra Sheet.

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
NONE

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EAST search terms:satellite, gateway, interface, client, server, packet\$4, protocol, (TCP or Internet), window size, round trip, delay, bandwidth

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 5,852,721 A (DILLON ET AL) 22 December 1998, see abstract, col. 4, lines 20-55, col. 7, line 22 to col. 8, line 2, col. 10, line 1 to col. 11, line 50, figures 1 and 15.	1-13, 17-44 ----- 14-16, 45-47
X,P ----- Y,P	US 5,968,129 A (DILLON ET AL) 19 October 1999, see abstract, col. 4, lines 14-54, col. 7, lines 14-61, col. 9, line 61 to col. 11, line 39, figures 1 and 15.	1-13, 17-44 ----- 14-16, 45-47, 53, 54, 60, 61, 67, 68
X,E ----- Y,E	US 6,038,216 A (PACKER) 14 March 2000, see col. 3, line 53 to col. 5, line 16, col. 7, line 66 to col. 8, line 67, figures 1, 2D-2H.	48-52, 55-59, 62-66, 69-99 ----- 53, 54, 60, 61, 67, 68

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"B" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

19 APRIL 2000

Date of mailing of the international search report

14 JUN 2000

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

ALPUS H. HSU

Telephone No. (703) 305-4377

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/02891

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,594,490 A (DAWSON ET AL) 14 January 1997, see entire document.	1-47
A	US 5,850,517 A (VERKLER ET AL) 15 December 1998, see entire document.	1-47
A,P	US 5,896,558 A (WIEDEMAN) 20 April 1999, see entire document.	1-47
A	US 5,163,046 A (HAHNE ET AL) 10 November 1992, see entire document.	48-99
A	US 5,426,635 A (MITRA ET AL) 20 June 1995, see entire document.	48-99

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/02891

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Please See Extra Sheet.

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
☒ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/02891

A. CLASSIFICATION OF SUBJECT MATTER:

US CL :

U.S. CL.: 370/235, 252, 349, 355, 401, 466; 455/12.1; 709/203, 218, 233

B. FIELDS SEARCHED

Minimum documentation searched

Classification System: U.S.

U.S. CL.: 370/229, 230, 231, 232, 235, 252, 253, 349, 352, 355, 389, 401, 465, 466, 467; 455/12.1, 13.1; 709/201, 203, 217, 218, 219, 230, 232, 233, 238, 249

BOX II. OBSERVATIONS WHERE UNITY OF INVENTION WAS LACKING

This ISA found multiple inventions as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single inventive concept under PCT Rule 13.1. In order for all inventions to be searched, the appropriate additional search fees must be paid.

Group I, claim(s) 1-47, drawn to a method, apparatus and system for providing transport of packetized information between client and server over a satellite network utilizing satellite gateways.

Group II, claim(s) 48-68, drawn to a method and system for managing memory for buffering packetized information in client-server communication system.

Group III, claim(s) 69-99, drawn to a method and system for controlling flow of packetized information over satellite in client-server communication system.

The inventions listed as Groups I-III do not relate to a single inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: Group I includes special technical feature of generic communication in client-server communication system over satellite. Group II includes special technical feature of memory management in client-server communication system. Group III includes special technical feature of data flow control in client-server communication system.